



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)


(51) International Patent Classification ⁶ : G06F 17/30	A1	(11) International Publication Number: WO 99/44157 (43) International Publication Date: 2 September 1999 (02.09.99)
<p>(21) International Application Number: PCT/US99/04146</p> <p>(22) International Filing Date: 25 February 1999 (25.02.99)</p> <p>(30) Priority Data: 60/076,048 26 February 1998 (26.02.98) US 09/044,835 20 March 1998 (20.03.98) US</p> <p>(71) Applicant: SUN MICROSYSTEMS, INC. [US/US]; 901 San Antonio Road, MS UPAL01-521, Palo Alto, CA 94303 (US).</p> <p>(72) Inventors: SCHEIFLER, Robert; 96 North Street, Somerville, MA 02144 (US). ARNOLD, Kenneth, C., R., C.; 7 Moon Hill Road, Lexington, MA 02173 (US). WALDO, James, H.; 155 Ruby Road, Dracut, MA 01826 (US).</p> <p>(74) Agents: GARRETT, Arthur, S.; Finnegan, Henderson, Farabow, Garrett & Dunner, L.L.P., 1300 I Street, N.W., Washington, DC 20005-3315 (US) et al.</p>	<p>(81) Designated States: AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, UZ, VN, YU, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SL, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).</p> <p>Published <i>With international search report.</i> <i>Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i></p>	

(54) Title: METHOD AND SYSTEM FOR MULTI-ENTRY AND MULTI-TEMPLATE MATCHING IN A DATABASE


(57) Abstract

A database system wherein one or more entry databases store a plurality of entries. Each entry is of a given type that defines the fields of the entry. Each field contains or identifies an object with associated attributes or data. The type of each entry may further define behavior in the form of methods the entry can implement. An entry type which is a sub-type of another inherits all fields and behavior of its super-type, and contains additional fields and/or defines new/modified behavior. Entries may be expressed in a Java™ programming language. The database system may further employ a search engine which allows queries to be made upon entries in the database. In one implementation, the queries include a read operation, a take operation, and a notify operation. Each query request includes a command indicating the type of operation, and a template which is an entry object having some or all of its fields set to specific values that must be matched exactly. Other fields of the template entry may be specified as wildcards, whose values do not matter. The search engine may further be configured to create an index for catching entries having characteristics which are commonly specified by queries. In one implementation, the databases may also store sets of entries which are matched using sets of templates.


ENTRY 1



ENTRY 2



ENTRY 3



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece			TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	NZ	New Zealand		
CM	Cameroon			PL	Poland		
CN	China	KR	Republic of Korea	PT	Portugal		
CU	Cuba	KZ	Kazakhstan	RO	Romania		
CZ	Czech Republic	LC	Saint Lucia	RU	Russian Federation		
DE	Germany	LI	Liechtenstein	SD	Sudan		
DK	Denmark	LK	Sri Lanka	SE	Sweden		
EE	Estonia	LR	Liberia	SG	Singapore		

METHOD AND SYSTEM FOR MULTI-ENTRY
AND MULTI-TEMPLATE MATCHING IN A DATABASE

RELATED APPLICATIONS

This application is a Continuation-In-Part of U.S. Application No. 08/971,529 filed on November 17, 1997, which is incorporated herein by reference.

5 The following identified U.S. patent applications are relied upon and are incorporated by reference in this application.

Provisional U.S. Patent Application No. 60/076,048, entitled "Distributed Computing System," filed on February 26, 1998.

10 U.S. Patent Application No. 09/044,923, entitled "Method and System for Leasing Storage," bearing attorney docket no. 06502.0011-01000, and filed on the same date herewith.

U.S. Patent Application No. 09/044,838, entitled "Method, Apparatus, and Product for Leasing of Delegation Certificates in a Distributed System," bearing attorney docket no. 06502.0011-02000, and filed on the same date herewith.

15 U.S. Patent Application No. 09/044,834, entitled "Method, Apparatus and Product for Leasing of Group Membership in a Distributed System," bearing attorney docket no. 06502.0011-03000, and filed on the same date herewith.

U.S. Patent Application No. 09/044,916, entitled "Leasing for Failure Detection," bearing attorney docket no. 06502.0011-04000, and filed on the same date herewith.

20 U.S. Patent Application No. 09/044,933, entitled "Method for Transporting Behavior in Event Based System," bearing attorney docket no. 06502.0054-00000, and filed on the same date herewith.

U.S. Patent Application No. 09/044,919, entitled "Deferred Reconstruction of Objects and Remote Loading for Event Notification in a Distributed System," bearing attorney docket no. 06502.0062-01000, and filed on the same date herewith.

25 U.S. Patent Application No. 09/044,938, entitled "Methods and Apparatus for Remote Method Invocation," bearing attorney docket no. 06502.0102-00000, and filed on the same date herewith.

U.S. Patent Application No. 09/045,652, entitled "Method and System for Deterministic Hashes to Identify Remote Methods," bearing attorney docket no. 06502.0103-00000, and filed on the same date herewith.

5 U.S. Patent Application No. 09/044,790, entitled "Method and Apparatus for Determining Status of Remote Objects in a Distributed System," bearing attorney docket no. 06502.0104-00000, and filed on the same date herewith.

U.S. Patent Application No. 09/044,930, entitled "Downloadable Smart Proxies for Performing Processing Associated with a Remote Procedure Call in a Distributed System," bearing attorney docket no. 06502.0105-00000, and filed on the same date herewith.

10 U.S. Patent Application No. 09/044,917, entitled "Suspension and Continuation of Remote Methods," bearing attorney docket no. 06502.0106-00000, and filed on the same date herewith.

15 U.S. Patent Application No. 09/044,839, entitled "Method and System for In-Place Modifications in a Database," bearing attorney docket no. 06502.0108, and filed on the same date herewith.

U.S. Patent Application No. 09/044,945, entitled "Method and System for Typesafe Attribute Matching in a Database," bearing attorney docket no. 06502.0109-00000, and filed on the same date herewith.

20 U.S. Patent Application No. 09/044,931, entitled "Dynamic Lookup Service in a Distributed System," bearing attorney docket no. 06502.0110-00000, and filed on the same date herewith.

U.S. Patent Application No. 09/044,939, entitled "Apparatus and Method for Providing Downloadable Code for Use in Communicating with a Device in a Distributed System," bearing attorney docket no. 06502.0112-00000, and filed on the same date herewith.

25 U.S. Patent Application No. 09/044,826, entitled "Method and System for Facilitating Access to a Lookup Service," bearing attorney docket no. 06502.0113-00000, and filed on the same date herewith.

30 U.S. Patent Application No. 09/044,932, entitled "Apparatus and Method for Dynamically Verifying Information in a Distributed System," bearing attorney docket no. 06502.0114-00000, and filed on the same date herewith.

U.S. Patent Application No. 09/030,840, entitled "Method and Apparatus for Dynamic Distributed Computing Over a Network," and filed on February 26, 1998.

U.S. Patent Application No. 09/044,936, entitled "An Interactive Design Tool for Persistent Shared Memory Spaces," bearing attorney docket no. 06502.0116-00000, and filed
5 on the same date herewith.

U.S. Patent Application No. 09/044,934, entitled "Polymorphic Token-Based Control," bearing attorney docket no. 06502.0117-00000, and filed on the same date herewith.

U.S. Patent Application No. 09/044,915, entitled "Stack-Based Access Control," bearing attorney docket no. 06502.0118-00000, and filed on the same date herewith.

10 U.S. Patent Application No. 09/044,944, entitled "Stack-Based Security Requirements," bearing attorney docket no. 06502.0119-00000, and filed on the same date herewith.

U.S. Patent Application No. 09/044,837, entitled "Per-Method Designation of Security Requirements," bearing attorney docket no. 06502.0120-00000, and filed on the same date
herewith.

15 BACKGROUND OF THE INVENTION

Field of the Invention

This invention relates to database systems and, more particularly, to entry storage in database systems. The invention also relates to query matching methodologies employed in database systems.

20 Description of the Related Art

Database systems are serving increasingly important roles in today's society. The amount of data that may be stored by typical database systems is growing at staggering rates, and providing expedient and flexible access to the data is often vital. However, the job of designing
25 database systems which perform in an optimal manner is difficult.

A particularly important consideration in database design is to store the data so that it can be used for a wide variety of applications and such that the manner in which it may be used can be changed quickly and easily. It is further often desirable to allow additional data types or categories to be added without significantly affecting other portions of the database system.

Unfortunately, it is remarkably difficult to change the way data is used and organized in many database systems, and to allow for additional data types and categories. Modifications can set off a chain reaction of changes to programs and other components of the database system, and hence can be exceedingly expensive to accomplish. Consequently, data processing has tended to become frozen into its old data structures. Additionally, designers of database systems tend to make numerous intricate trade-offs between one aspect of the design and another, thus leading to a less than optimal resulting design.

To elaborate further, storing information in a database system is usually done by exact type; that is, a record stored in a database system typically has a certain number and type of fields, such as a tuple in a relation of a relational database. Each field may be defined to store a particular type of data, such as string data, integer data, boolean data, or some other data type, as desired. If a user wants to look-up information in the database, the user must look-up the exact record type (i.e., records of different types are typically searched independently). Thus, if data in which the searcher is interested may be in a record of any one of several different types, the searcher must know in advance which record types might exist, or have some explicit way of asking for all types. The searcher is then required to explicitly look through all possible record types to find the particular record of interest. This complicates the software associated with the search engine and can practically limit the number of different record types that can be maintained.

Furthermore, fields in a record in a database typically do not refer to other records. This reduces the amount and kind of information a single record typically holds. Consequently, databases are limited by their inability to store and manipulate records that are sets of records.

In addition, a record in a typical database is usually not associated with characteristics apart from the actual data itself. For example, when data is found, there is no way to associate behavior with the record that changes with the record type, thus allowing the program to use various records returned by the match in a single way to achieve record-type specific results.

Regarding another aspect of database systems, the usual way of specifying attributes or properties in a database is by a simple set of <name, value> pairs, where the values are either restricted to strings or else are of arbitrary type (type "any"). There is no way to restrict type in these <name, value> pairs, and hence, there is no enforcement that a particular name must

always have a value of a particular type. For example, in a hypothetical printer, the "speed" attribute, representing how many pages-per-minute the printer can print, should be enforced as an integer to ensure a user does not set it to something else, such as "about 10 ppm" or "8 to 10." The same kind of problem exists if the values can be objects of any type. In that case, a user can set the attribute to an integer object instead of a string object but still cannot prevent another user from incorrectly storing a string object in an attribute that should only be set to integers.

Furthermore, previous systems have no enforcement of relationships between different attributes (*e.g.*, that one attribute should exist if and only if another attribute also exists). For example, given three exemplary attributes for a printer: resolution (dots per inch), speed (pages per minute), and color (vs. black and white), in reality, many printers can perform different combinations of these attributes. For example, a typical color printer might be described as:

5 pages per minute at 600 dots per inch in black and white;

7 pages per minute at 300 dots per inch in black and white;

1 page per minute at 300 dots per inch in color.

Whereas a given attribute can have multiple values, some conventional systems allow attributes to have multiple values, and these systems might represent the given attributes in the following manner:

speed = {1, 5, 7}

resolution = {300, 600}

color = {true, false}

but this does not express the interrelationships between the pairs of attributes. Expressly, there is no representation and enforcement of the three interrelationships:

{speed: 1, resolution: 300, color: true};

{speed: 7, resolution: 300, color: false};

{speed: 5, resolution: 600, color: false}.

Accordingly, it is desirable to provide a database system allowing for polymorphic entry, entry matching and manipulation including sets of entries, to thereby allow more efficient searching and to allow simplified software implementation where many different types exist. It would further be desirable to provide a system where entries may define their own behavior

according to type. Additionally, it would be desirable to have type-safe attribute matching in a database system and have enforcement of relationship between attributes.

Summary of the Invention

5 The problems outlined above are solved by a method and system for storing and searching entries in a database consistent with the present invention. In one embodiment, a computer system implements a database system wherein one or more entry databases store a plurality of entries. Each entry is of a given type that defines the fields of the entry. Each of the fields contains or identifies an object with associated attributes or data. The type of each entry
10 may further define behavior in the form of methods the entry is configured to implement. An entry type which is a subtype of another inherits all of the fields and behavior of its supertype, and contains additional fields and/or defines new or modified behavior. Entries may be expressed in a Java™ class of the Java™ programming language. Similarly, in one implementation, each field is expressed in a defined class.

15 The database system may further employ a search engine which allows queries to be made upon entries in the database. In one implementation, the queries include a read operation, a take operation, and a notify operation. Each query request includes a command indicating the type of operation and a template which is an entry object having some or all of its fields set to specified values that must be matched exactly. Other fields of the template entry may be
20 specified as wildcards, whose values do not matter. The polymorphic entry and entry matching technique advantageously allows a query to be made which might result in a match of an entry which is the same type as the template entry or which is a subtype of the template entry. This allows for efficient searching and for simplified implementation of the search engine. Furthermore, the search engine need not be aware of all possible subtypes which might exist.
25 Finally, since each entry has an associated class, an entry may define its own behavior by supporting one or more methods. A method of a particular entry may be invoked by a program to implement various functionality, such as equivalence checking.

5 The search engine may further be configured to create an index for entries having characteristics which are commonly specified by queries. That is, the search engine may detect a pattern of usage and create an index which points to entries having attributes which are commonly specified by received templates. In this manner, expedient entry searching can be accomplished.

10 Additionally, the search engine may further be configured to store sets of entries in a single record. These sets of entries are matched using sets of templates. In one implementation, operations similar to single entry implementation operations, such as read, take, write, and notify, may be performed on sets of entries.

Brief Description of the Drawings

Fig. 1 is a diagram of a distributed computer system employing a database system consistent with the present invention;

15 Fig. 2 is a diagram illustrating exemplary communication between a plurality of clients and a plurality of entry databases consistent with the present invention;

Fig. 3 is a diagram illustrating several entries each including fields of particular types consistent with the present invention. Fig. 3 further illustrates an entry which is a subtype of another entry;

20 Fig. 4 is a diagram illustrating various operations supported by a database system consistent with the present invention;

Fig. 5 is a diagram illustrating an exemplary template entry consistent with the present invention;

25 Fig. 6 is a flow diagram which depicts one implementation of a method for searching entries within a particular entry database in response to either a read request or a take request consistent with the present invention;

Fig. 7 is a functional block diagram illustrating aspects of a database system consistent with the present invention;

Fig. 8 is a flow diagram which illustrates an implementation of a search methodology to process a notify request consistent with the present invention;

Fig. 9A and 9B depict an implementation of an indexing methodology consistent with the present invention;

Fig. 10 is a diagram illustrating several multi-entries, each multi-entry including fields containing entries;

5 Fig. 11 is a diagram illustrating various multi-entry operations supported by a database system consistent with the present invention;

Fig. 12 is a diagram illustrating an exemplary multi-template consistent with the present invention;

10 Fig. 13 is a flowchart depicting a method of searching a database for multi-entries consistent with the present invention;

Fig. 14 is a functional diagram illustrating aspects of a database system in relation to an example of a multi-entry search consistent with the present invention;

Fig. 15 is a flowchart which illustrates an implementation consistent with the present invention of a search methodology to process a notify request for multi-entries;

15 Fig. 16 is a diagram illustrating various operations to modify entries and multi-entries in place in the database consistent with the present invention;

Fig. 17 is a flowchart illustrating steps of a single entry modify operation consistent with the present invention;

20 Fig. 18 is a flowchart illustrating steps of a multi-entry add operation consistent with the present invention;

Fig. 19 is a flowchart illustrating steps of a multi-entry modify/delete operation consistent with the present invention; and

Fig. 20 is a flowchart illustrating steps of a notification operation for use with in-place modification operations consistent with the present invention.

25 While the invention is susceptible to various modifications and alternative forms, specific embodiments thereof are shown by way of example in the drawings and will herein be described in detail. It should be understood, however, that the drawings and detailed description thereto are not intended to limit the invention to the particular form disclosed, but on the contrary, the intention is to cover all modifications, equivalents and alternatives falling within the spirit and
30 scope of the present invention as defined by the appended claims.

Detailed Description of the Invention

Methods and systems consistent with the present invention store and search entries in a database. Each entry is of a given type which defines the fields of the type, and each field contains or identifies an object of a given type. These entries may be expressed in a Java™ class of the Java™ programming language. Additionally, methods and systems consistent with the present invention may also store and search sets of entries in a database. Generally, these searches are performed using sets of templates.

Distributed System

Turning now to the drawings, Fig. 1 is a diagram illustrating a distributed system 100 in which polymorphic entry, entry matching and entry manipulation in consistent with the present invention may be employed. Fig. 1 shows a plurality of computer systems 101-103 interconnected by a network 105. Network 105 is illustrative of, for example, a local area network (LAN), a wide area network (WAN), or the Internet, among a variety of other possible communication channels. Each computer system includes a memory and a microprocessor such as the exemplary memory 120 and CPU 122 of computer system 101. Memory 120 is representative of various types of possible memory, such as hard disk storage, floppy storage, and dynamic random access memory (DRAM). Installation media 126, which is illustrative of one or more CD ROMs or floppy disks (or other storage media), is further shown in association with computer system 101. Generally speaking, installation media 126 may store code and/or data which implements portions of the search engine and other aspects of the database system described below and may be utilized to install the code and/or data within a particular computer system. It is noted that in other embodiments, a database system in accordance with the invention may be implemented within a single, non-networked computer system.

As will be described in further detail below, computer system 100 implements a database system wherein object oriented techniques are utilized to store entries of possible types or subtypes. This database may also store sets of entries, called multi-entries, as described after the single entry implementation. Entries of the database may be stored solely within memory 122 of computer system 101 or may be distributed across the memories of several different systems, such as memory 122 and similar memories of computer systems 102 and 103. Each entry is of a given type that defines the fields of the entry. Each of the fields contains or identifies an object

with associated attributes or data. The type of each entry may further define behavior in the form of methods the entry is configured to implement. An entry type which is a subtype of another inherits all of the fields and behavior of its supertype and contains additional fields and/or defines new/modified behavior.

5 The database system further employs a search engine, which may be implemented in software, which allows queries to be made upon entries (or multi-entries, described later) in the database. In one implementation, each query includes a template that specifies the values (or attributes) of the key fields of a given entry which must match corresponding fields of the template to result in a successful match. Other fields of the template entry may be specified as
10 wildcards, whose values do not matter. The polymorphic entry and entry matching technique employed by computer system 100 advantageously allows a query to be made which might result in a match of an entry which is the same type as the template entry or which is a subtype of the template entry. Additional details regarding aspects of the polymorphic entry and entry matching technique are provided further below in conjunction with Figs. 2-9.

15 Fig. 2 is a diagram illustrating exemplary communication flow between a pair of clients (or users) 202 and 204 and a plurality of entry databases 206-208. Each of the entry databases 206-208 may be distributed across one or more memory subsystems of distributed system 100. Each client 202 and 204 is configured to generate queries which may be provided to one or more of the entry databases 206-208. Clients 202 and 204 may further be configured to store (write)
20 entries within entry databases 206-208, as discussed further below. Fig. 2 finally illustrates an event catcher 210, the functionality of which will also be described further below.

Entry Matching and Operation

Fig. 3 is a diagram illustrating a plurality of exemplary entries which may be stored in any of entry databases 206-208. In one implementation, each entry is a typed group of objects expressed in a Java™ class of the Java™ programming language. The Java™ programming
25 language is described in Ken Arnold, James Gosling, The Java™ Programming Language, Addison-Wesley, 1996, which is hereby incorporated by reference. Java™ is a registered trademark of Sun Microsystems, Inc. In the United States and other countries. It is understood, however, that in other embodiments, entries may be defined in classes of other programming
30 languages.

Entry 1 of Fig. 3 illustratively includes a set of four fields (field A, field B, field C and field D) which are defined by the class type of entry 1. Entry 2 includes a set of five fields (field A, field B, field C, field D and field E) which are defined by the class type of entry 2. It is noted that entry 2 is a subtype of the type of entry 1, since entry 2 contains the same field types as entry 1, in addition to others.

Entry 3 includes a set of four fields (field W, field X, field Y and field Z) defined by the class type of entry 3. It is noted that while entries 1 and 3 each contain four fields, they represent different types since their field types are not identical.

Numerous other types and subtypes may further be associated with the database system, as desired depending upon the requirements of a particular application or database. Additionally, it is noted that entries are possible which are subtypes of subtypes. For purposes of example only, the following code is illustrative of specific exemplary entries defined in the Java™ programming language:

A "Request For Bid" object for book buying:

```
public class RFB implements Entry {  
    public Integer    order number; \\ The buyer's tracking number for  
                                   this order.  
    public String     title;        \\ The title of the book.  
    public Integer    count;        \\ The number of copies desired.  
}
```

A "Full" request for bids, that is, one that includes the ISBN:

```
public class FullRFB extends a RFB {  
    public String     isbn; \\ The ISBN field.  
}
```

Referring again to Fig. 3, each of the fields of the entries are configured to store values which point to or identify typed objects. Each of these typed objects may similarly be expressed in a Java™ class. For example, field A of entry 1 contains a value identifying objects 310, and field B of entry 1 identifies an object 312. Similarly, field A of entry 2 identifies an object 314,

and field B identifies an object 316, and so on. It is noted that for the exemplary entries of Fig. 3, since entry 2 is a subtype of entry 1, the object 310 associated with field A of entry 1 is expressed by the same class type as that of the object 314 associated with field A of entry 2. The attributes or data associated with each of the objects 310 and 314, however, may differ. It is further noted that the objects identified by the fields of the various entries may contain integer data, boolean data, string data, or other data types, as desired, depending upon the defined type. Still further, it is noted that since each field is associated with a typed object, methods may be supported and selectively invoked via external stimuli to perform operations upon the data, as desired. Similarly, each entry type may support one or more methods, as described further below.

Referring back to Fig. 2, as described previously, each of the entry databases 206, 207 and 208 may store numerous entries of various types of subtypes. An entry may be stored within a given entry database 206-208 through a write operation invoked by a client. Entries may additionally be provided to an entry database via an installation media (e.g., installation media 126 of Fig. 1).

Several additional operations are further supported by the system, as illustrated by Fig. 4. Specifically in addition to write operations, a client or user may initiate "read" operations, "take" operations, and "notify" operations.

Read, take, and notify operations are each used to query a particular entry database to determine whether an entry with certain specified characteristics has been written into an entry database and, if so, to cause a specific action. In one particular implementation, these queries are made through the use of templates included as part of the operation requests which define the type of entries which are of interest with respect to the operation. A read operation may be initiated to read an entry in an entry database which matches an associated template. A take operation is similar to a read; however, a take operation will remove an entry matching the template from the entry database. Finally, a notify operation registers interest in future incoming entries of a specified entry database that match the associated template. When matching entries arrive, the client or a specified event catcher is notified. Further details regarding each of these operations will be provided further below.

Fig. 5 depicts an exemplary template. Generally speaking, a template specifies a particular entry type whose fields can either have values (reference to objects) or "wildcards." When considering a particular entry in an entry database as a potential match against a template, the attributes of specified fields (i.e., those which do not correspond to wildcarded fields of the template) of the entry must exactly match the attributes (or fall within a range of possible attributes) specified by the corresponding field of the template. Fields of an entry which correspond to "wildcard" fields of the template may have any attributes and do not affect the matching outcome (i.e., wildcards in the template match any value in the same field of an entry). In one implementation, a wildcard is specified with a null reference value. As will be described further below, the type of an entry which matches a template may be a subtype of the type of the template. Thus, all fields of a particular entry which are defined by the subtype are considered to be wildcards. This enables a template to match entries of any subtype. It is noted that an entry having no wildcards may further be a valid template.

Fig. 6 is a flowchart which depicts one implementation of a method for searching entries within a particular entry database in response to either a read request or a take request consistent with the present invention. As stated previously, a read request may be initiated by a client (or user) to read an entry that matches a given template. A take request may be initiated to remove an entry matching the given template from an entry database.

(1) Method for Entry Matching

The steps illustrated in Fig. 6 will be described in conjunction with Fig. 7, which depicts functional aspects of a database system employing the search methodology of Fig. 6. The functional components of Fig. 7 may be implemented through a combination of software and hardware associated with one or more of the computer systems 101-103 of Fig. 1. During step 602, a read or take request is provided from a user 702 to a database manager 704. A given request will include a specified template, an exemplary one of which is shown as template 706 in Fig. 7. A search engine 708 of database manager 704 subsequently searches entry database 710 during step 604 to determine if any matching types or subtypes of the template reside within entry database 710. If no entries within entry databases 710 are of the same type or a subtype of the template (step 606), database manager 704 returns a null value to user 702 during step 608.

If the search engine 708 determines that there are one or more entries within entry database 710 that are matching types or subtypes of the template, it compares the fields of a selected entry (of the same type or a subtype) with the corresponding fields of template 706 which are not designated as wildcards (step 610). The search engine 708 disregards any fields of the selected entry which correspond to a wildcard field of the template (step 612). In one specific implementation, the search engine 708 determines whether the attributes or data associated with each non-wildcarded field of an entry exactly matches the attributes or data of the corresponding template entries. If the attributes of any field which is not a wildcard field do not exactly match those of the template, it is determined that a match with respect to that entry has not occurred. It is noted that in other embodiments, a field of the template may specify a range of values (rather than requiring an exact match) in which the data of a corresponding field must fall to result in a match.

(2) Example

In the example illustrated by Fig. 7, field C of template 706 is specified as a wildcard. Thus, when the search engine 708 searches entry database 710 for a matching entry, the search engine 708 first determines whether any entries reside within the entry database 710 which are types or subtypes of the template 706. As shown, both entry 1 and entry 2 meet this qualification. The search engine 708 then selects one of these entries (such as entry 2, as depicted) and compares the fields of the entry to the fields of the template which are not specified as wildcards. The search engine 708 determines whether an exact match for each of the non-wildcard fields of the template occurs with respect to the attributes of each of the corresponding fields of the entry. It is noted that since field C of the template 706 of Fig. 7 is designated as a wildcard field, the attributes of field C of entry 2 are disregarded and do not affect the comparison operation. Similarly, those fields of any entry defined by a subtype of the template type are disregarded (i.e., a match may still occur even though the attributes of the wildcard fields do not match). If a matching entry is found (all non-wildcarded fields match), the search engine 708 either reads the matching entry from the entry database 710 and provides it to the user, or removes the entry from entry database 710, depending upon whether the request was a read request or a take request. The search engine 708 searches entry database 710 until either a match is found or until all of the entries have been checked. If no matching entries are

found after a search of all entries in the database, database manager 704 returns a null value (step 616) to user 702.

(3) Notify

Turning next to Fig. 8, a flow diagram is shown which illustrates one implementation of a methodology implemented by the search engine 708 to process a notify request received from a user consistent with the present invention. As stated previously, a notify request may be issued by the user to register interest in future incoming entries which match a template. When matching entries arrive, the database manager either notifies the user or an event catcher (Fig. 2), as desired. When a user invokes a notify request, a lease time is specified which indicates how long the database manager 704 will register the notification. After the specified time has elapsed, the notify registration is canceled.

Thus, as depicted in Fig. 8, when database manager 704 receives a new notify request (step 802), incoming entries are monitored until the registration time expires (steps 804 and 806). If an entry is written into entry database 710 which is a type or a subtype of the template (step 808), the fields of the entry are compared to determine whether their attributes match the attributes associated with the non-wildcard fields of the template (step 810). Similar to the previous discussion, fields of an incoming entry which are wildcarded or which are defined by a subtype of the template entry are disregarded (step 812). If a match occurs, the user or an event catcher is notified (steps 814 and 816). Additional incoming entries continue to be monitored as long as the registration time has not expired.

(4) Indexing

Turning back to Fig. 7, the search engine 708 may further be configured to create an index for entries having characteristics which are commonly specified by queries received by database manager 704. That is, the search engine 708 may detect a pattern of usage and create an index 720 which points to entries having attributes which are commonly specified by received templates. As an example, consider a situation in which database manager 704 receives several query requests each using templates which are of the same type or a subtype of the template 706 illustrated in Fig. 7. In this situation, the search engine 708 may detect that the same attribute value for field A (such as an integer value 1,000), is specified in each of the templates of the numerous requests. In one implementation, database manager 704 is configured to detect such

a common template usage specified by many queries and to responsively create an index 720 which points to all entries within entry database 710 which are the same type or a subtype of the commonly received template and which all have the common attribute specified for field A. Consequently, when a new query with a similar template is received (that is, a template which again specifies the value 1,000 for field A), the search engine 708 may use index 720 to quickly locate the subset of entries which might match the template, without searching through the entire set of entries in the database.

An implementation of the indexing methodology is illustrated in Figs 9A and 9B. As illustrated in Fig. 9A, an index to entries corresponding to a particular template category may be created by tracking request templates (step 910) and determining if there is a pattern associated with the templates (e.g., templates in which a particular non-wildcarded field (or fields) is specified with the same attribute or value (step 912). This determination may be made, for example, based on receiving a predetermined number of queries within a particular time which have similar templates, or based upon receiving a certain percentage of queries which have similar templates. When the search engine 708 detects common template usage pattern, it creates an index to the entries which are of the particular template category (step 914). Thus, considering the above example, the search engine 708 may create an index to all entries of the same type or subtype as template 706 with field A value of 1,000.

As depicted in Fig. 9B, if a new request is received during step 920, a normal search is performed if an index does not exist for that particular template category (steps 922 and 924). If an index category exists for the template of the new request, the index is used to find the subset of potentially matching entries and compares only those entries. In this manner, expedient entry searching may be accommodated. It is noted that the search engine 708 may subsequently delete certain indexes if the common usage is no longer detected, as desired.

As stated previously, an entry type may further define a method which may be invoked by external stimuli (e.g., a software program associated with the database system that processes the entries) to cause a specific action to occur. Such specific actions might affect data residing in the entry. In addition, a method may be utilized to allow efficient equivalence checking to determine if a particular entry is equivalent to another (or equivalent to a template), even though the form of the data may be different. For example, an entry that stores decompressed video data

might be considered an equivalent to another entry storing a compressed form of the video data. A method may be provided to determine such entry equivalence, even though the actual stored data is different (e.g., a method may be invoked to determine whether the objects refer to the same underlying pathname in the file system associated with the image).

5 It is further noted that entries may be stored in an implementation-specific representation that includes a serialized form of an entry's class and each of the entry's fields. Serialized forms of a template may then be compared against the serialized forms of the entries. It is finally noted that an entry may be specified by the field of another entry.

Multi-Entry and Multi-Template Matching

10 An alternative embodiment implements operations and matching of items in the database, each item defining a set of entries. An "item" is referred to herein as a "multi-entry"; "multi-template" (or "item template") is the term used to refer to a set of templates, which are used to match items in the database. This section will further describe the matching of multi-entries to multi-templates and operations involving multi-entries and multi-templates. Much of the
15 behavior of multi-entries and multi-templates is similar to single entry and template behavior as discussed previously.

Fig. 10 is a diagram illustrating a plurality of exemplary multi-entries which may be stored in any of the entry databases 206-208. In methods and systems consistent with the present invention, each multi-entry is a group of entries, each entry being a typed group of objects
20 expressed, for example, in a Java™ class of the Java™ programming language. In other words, a multi-entry is set of entries and has no type of its own.

Referring to Fig. 10, multi-entry 1 illustratively includes a set of four elements (element A, element B, element C and element D). Whereas multi-entries may be of varying length, the exemplary multi-entry 2 includes a set of five elements (element E, element F, element G,
25 element H and element I), and the exemplary multi-entry 3 includes a set of three elements (element J, element K, and element L). Each of the elements of the multi-entries is configured to store typed entries. As stated, each of these typed entries may be expressed in a Java™ class, and the fields of the entries store values which identify typed objects. For example, element A of multi-entry 1 contains a value identifying entry 1010, and element B of multi-entry 1 identifies
30 an entry 1012. Similarly, element E of multi-entry 2 identifies an entry 1014, and element F

identifies an entry 1016, and so on. It is noted that for the exemplary multi-entries of Fig. 10, while any entries may be of the same type, the attributes or data associated with the fields of the entries, however, may differ. It is further noted that although the entries identified by the elements of various multi-entries have a specified type, the multi-entries are simply sets of entries and themselves have no specific type.

Since multi-entries and multi-templates are just sets, they can be represented using any suitable representation of a set. For purposes of example only, a suitable representation in many programming languages, including the Java programming language, is an array, where the size of the array is given by the number of entries in the multi-entry or the number of templates in the multi-template, and each element of the array is an entry (in the case of multi-entries) or a template (in the case of multi-templates).

Referring back to Fig. 2, as described previously, each of the entry databases 206, 207 and 208 may store numerous multi-entries. A multi-entry may be stored within a given entry database 206-208 through a write operation invoked by a client. Multi-entries may additionally be provided to an entry database via an installation media (e.g., installation media 126 of Fig. 1). Several additional operations, similar to single entry operations, are further supported by the system, as illustrated by Fig. 11. Specifically in addition to write operations, a client or user may initiate "read" operations, "take" operations, and "notify" operations.

Read, take, and notify operations operate on multi-entries in the same manner as they operate on single entries. In one particular implementation, these queries are made through the use of multi-templates included as part of the operation requests which define which multi-entries are of interest with respect to the operation. As in the single entry implementation, the read operation reads multi-entries in the database that match the given multi-template. Once again, the take operation performs the same function as the read operation but also removes any multi-entry read from the database. Used to register interest in future multi-entries, the notify operation notifies the client or an event catcher when matching multi-entries arrive in the database.

Fig. 12 depicts an exemplary multi-template. Generally speaking, a multi-template specifies a particular multi-entry whose elements contain templates as previously defined. The multi-template 1200 shown in Fig. 12 has four elements, W, X, Y and Z, although it should be

noted that multi-templates can have any number of elements, but, in an exemplary embodiment of the present invention, each element in a multi-template 1200 includes one template 1202.

(1) Method for Multi-Entry and Multi-Template Matching

Generally, a multi-template will match a multi-entry if each template in the multi-
5 template has a matching entry in the multi-entry. Each entry in the multi-entry can be used to match more than one template in the multi-template. For example, a single entry in the multi-entry could match all the templates in the multi-template, thus creating a match between the multi-entry and the multi-template. In the process of determining whether a multi-entry matches a multi-template, it is necessary to compare the individual entries of the multi-entry with the
10 individual templates in the multi-template. This matching of individual templates in a multi-template to individual entries in a multi-entry is performed in the same manner as the matching of single entries to single templates as previously defined.

Fig. 13 is a flowchart which depicts one implementation of a method for searching multi-entries within a particular entry database in response to either a read request or a take request.
15 The steps illustrated in Fig. 13 will be described in conjunction with Fig. 14, which depicts functional aspects of a database system employing the search methodology of Fig. 13. The functional components of Fig. 14 may be implemented through a combination of software and hardware associated with one or more of the computer systems 101-103 of Fig. 1. During step 1302, a read or take request is provided from a user 702 to a database manager 704. A given
20 request will include a specified multi-template, an exemplary one of which is shown as multi-template 1200 in Fig. 14. A search engine 708 of database manager 704 subsequently searches entry database 710 to determine if any matching multi-entries of the multi-template reside within entry database 710.

To search the entry database 710, the search engine 708 selects a multi-entry in the entry
25 database 710 to compare to the given multi-template (step 1304). Then, to determine if the multi-entry matches the specified multi-template, the search engine 708 compares each template in the multi-template to the entries of the multi-entry (step 1306). If there is a matching entry for each and every template in the multi-template, the multi-entry matches the multi-template (step 1308). Each entry in the multi-entry can be used to match more than one template in the
30 multi-template. If it is determined that every template in the multi-template has a matching entry

in the multi-entry (step 1310), the multi-template matches the multi-entry, and the multi-entry is returned (step 1312). This process continues until all entries in the database have been examined.

5 If a match is not made between the selected multi-entry and the multi-template supplied by the operation, the search engine 708 searches other multi-entries in the entry database 710 (step 1314). If no multi-entries are found in the database to match the multi-template, a null value is returned (step 1316). If more than one is found, in one implementation, one multi-entry is returned, and in another implementation, a requested number of matching multi-entries are returned. If the original request from the client arrived in the form of a take request, the
10 matching multi-entry is removed from the database. However, if the original request arrived as a read request, the matching multi-entry is not removed from the database.

(2) Example

In the example illustrated by Fig. 14, the search engine 708 selects one of the multi-entries (such as multi-entry 2, as depicted) and compares the entries of the multi-entry to the
15 templates of the multi-template. In this example, the search engine 708 receives a request including the multi-template 1200.

The search engine 708 determines whether the templates in the multi-template 1200 match the entries of the selected multi-entry in the entry database 710. The search engine 708 compares the template in element W of the multi-template 1200 to the entries in elements E, F,
20 G and H and I of the multi-entry 2 in the entry database 710 to determine whether any of these entries match the template. If the search engine 708 finds a match for the template in element W of multi-template 1200, it can then attempt to find a match for the template in element X of multi-template 1200. It will compare the template in element X of the multi-template 1200, once again, to each of the entries in the multi-entry 2 until it finds a match. It should be noted that
25 even though each of these entries in elements E, F, G, H and I may have been compared to other templates in the multi-template (*i.e.*, the template in element W), they can be compared again to this template. If a match is found for the template in element X, the search engine 708 checks the rest of the templates in the multi-template 1200 (the templates in elements Y and Z) in the same manner. Since each entry can be compared to more than one template, the entry in element
30 E of multi-entry 2 could, for example, match all of the templates in the elements W, X, Y and

Z of the multi-template 1200. If each template in the multi-template 1200 has a matching entry in multi-entry 2, the multi-entry 2 matches the multi-template 1200.

As with single entry operations, if a matching multi-entry is found, the search engine 708 either reads or removes the matching multi-entry from the entry database 710 and provides it to the user, depending upon whether the request was a read request or a take request. The search engine 708 continues searching the entry database 710 until either a match is found or all of the multi-entries have been checked. If no matching multi-entries are found after a search of all multi-entries in the entry database 710, the database manager 704 returns a null value to user 702.

(3) Notification

Turning next to Fig. 15, a flowchart is shown which illustrates one implementation of a methodology implemented by search engine 708 to process a notify request received from a user. The notification operation works similarly to the notification for single entries, and as stated previously, a notify request may be issued by the user to register interest in future incoming multi-entries which match a multi-template. When matching multi-entries arrive, the database manager 704 either notifies the user or an event catcher (Fig. 2), as desired. When a user invokes a notify request, a lease time is specified which indicates how long the database manager 704 will register the notification. After the specified time has elapsed, the notify registration is canceled.

Thus, as depicted in Fig. 15, when database manager 704 receives a new notify request (step 1502), incoming multi-entries are monitored until the registration time expires (steps 1504 and 1506). If a multi-entry is written into entry database 710, the search engine 708 determines whether the multi-entry matches the multi-template. Similar to the previous discussion, if each template of the multi-template has a matching entry in the multi-entry (steps 1508 and 1510), a match occurs and the user or the event catcher is notified (step 1512). Additional incoming multi-entries continue to be monitored as long as the registration time has not expired (step 1504).

In-Place Entry Modifications

An alternative embodiment additionally implements modifications of entries and multi-entries in place in the database. Consistent with the present invention, this implementation is

an extension to the entry/template and multi-entry/multi-template scenario in which the fields of entries can be modified in place in the database, and in the multi-entry case, entries within a multi-entry can be added, modified or deleted. The term "in-place" means that the modification occurs in the database without removing the entries or multi-entries from the database.

5 In-place modification operations consistent with this implementation of the present invention include a "single entry modify," a "multi-entry add," a "multi-entry modified/delete," and a "notify." Fig. 16 illustrates these in-place modifications and the information supplied by a user to request a modification. The single entry modify operation changes the values in the individual fields of a entries in the database. The multi-entry add operation adds entries to multi-
10 entries existing in the database, and the multi-entry modified/delete changes or deletes entries in multi-entries in the database. The notify operation, with respect to in-place modifications, notifies of a match occurring due to a change in an entry or multi-entry in the database. Generally, the first parameter of an operation request, as shown in Fig. 16, specifies which entries or multi-entries in the database are to be modified, and the second parameter specifies
15 the modifications to be made.

(1) Single Entry Modify

Fig. 17 depicts a flowchart illustrating an implementation of a single entry modify operation consistent with the present invention. Further reference will also be made back to Fig. 7 for structural aspects. To invoke a single entry modify operation to change the values of
20 individual fields in an entry, a user 702 makes a request including two templates (step 1700, see also Fig. 16). The first template 1602 is used to identify the entries in the database to be modified while the second template 1604 is used to modify the entries and identifies the modifications to be made. Since these two templates operate on the same entry by identifying it and modifying it, the second template 1604 is the same type as, or a supertype of, the first
25 template 1602. In certain implementations, the single entry modify procedure may include a type checking of the templates.

First, the search engine 708 searches for entries in the entry database 710 matching the first template 1602 received in the request from the user 702 (step 1702). When comparing the entries in the database to the first template, the matching is done according to the single entry
30 and template matching as discussed above. If there are no entries in the entry database 710

matching the first template 1602 (step 1704), the search engine 708 returns a zero to the user 702 representing the number of entries changed due to this operation (step 1712).

If a matching entry is found, the second template 1604 is used to modify the entry in place in the database, *i.e.*, the entry is not removed from the database for modification. Because the second template 1604 is the same type as, or a supertype of, the first template 1602 and the first template 1602 is the same type as, or a supertype of, the entry being modified, therefore the second template 1604 is also the same type as, or a supertype of, the entry being modified, and so the second template 1604 has the same, or a subset of, fields of the entry being modified. To modify the matching entry, the values in the non-null fields of the second template 1604 are used to replace the values in the corresponding fields of the entry (step 1708). The value of the fields of the entry corresponding to the null or wildcard fields of the second template 1604 remain unchanged. After the modification of an entry in the database, the search engine 708 continues searching for more entries to be modified. If there are any more entries in database matching the first template 1602 (step 1710), those entries are modified in the same manner (step 1708). Upon completion, the search engine 708 returns a number representing the number of entries modified by the operation (step 1712).

(2) Multi-Entry Add

Fig. 18 is a flowchart illustrating steps of a multi-entry add operation consistent with the present invention. Generally, the multi-entry add operation adds entries to specified multi-entries in the database. A user 702 requesting this operation supplies two parameters to the search engine 708, a multi-template 1606 and a multi-entry 1608 (see Fig. 16). The first parameter, the multi-template 1606, is used to identify multi-entries in the entry database to be changed; the second parameter, the multi-entry 1608, is a set of entries to be added to the identified multi-entries database. Using multi-entry/multi-template matching as previously discussed, the search engine 708 searches for multi-entries in the database that match the multi-template 1606 given by the user 702 (step 1802). If no matching multi-entry is found (step 1804), the search engine 708 returns a zero to the user 702 representing the number of multi-entries to which entries were added (step 1812).

If a matching multi-entry is found (step 1804), the entries of the specified multi-entry 1608 supplied by the user 702 are added to the matching multi-entry found in the entry database 710 without the removal of the matching multi-entry from the entry database 710 (step 1808). After adding to a multi-entry in the database, the search engine 708 continues searching for more multi-entries to change. If there are more multi-entries in the entry database 710 that match the multi-template 1606 given as the first parameter in the operation request (step 1810), the entries of the multi-entry 1608 in the operation request are added to these multi-entries in the same manner (step 1808). Finally, the search engine 708 returns a number representing the number of multi-entries to which entries were added (step 1812).

(3) Multi-Entry Modify/Delete

Fig. 19 is a flowchart illustrating a method of implementing a multi-entry modify/delete operation consistent with the present invention. The multi-entry modify/delete operation either modifies entries in multi-entries in the database or deletes entries from multi-entries from the database, depending on the entries in the multi-entry in the request. When the search engine 708 receives a multi-entry modify/delete request, the request includes two multi-templates, a multi-template 1610 and a multi-template 1612, as its two parameters (step 1900, see also Fig. 16). The multi-template 1610 is used to identify the multi-entries in the entry database 710 that are to be modified or have entries deleted, while the multi-template 1612 supplied by the user 702 specifies the modification and deletions to be made to the entries in a matching multi-entry.

The multi-template 1612, given in the request, contains the same number of templates as the multi-template 1610. Furthermore, each individual template in the multi-template 1612 is the same type, or a supertype of, as the corresponding template in the multi-template 1610 (*i.e.* every template (i) in multi-template 1612 is the same type as, or a supertype of, template (i) in multi-template 1610). For instance, the first template in the multi-template 1612 is the same type as, or a supertype of, the first template in the multi-template 1610, the second template is the same type as, or a supertype of, the corresponding second template and so on. Differing from a common multi-template, the templates in the multi-template 1612 of the request parameter may also be null values, and these templates will be used to designate deletions from a multi-entry as described below.

The search engine 708 first searches for multi-entries matching the multi-template given as the first parameter in the operation request (step 1902). If no match is found, a zero is returned representing the number of multi-entries that were modified (step 1914).

5 If a matching multi-entry is found in the entry database 710, it is modified using the multi-template 1612 from the operation request. To accomplish this modification, for each non-null template in the multi-template 1612, the corresponding entry in the matching multi-entry in the entry database 710 is changed according to the non-null template in multi-template 1612 (step 1908). This change is performed by using single entry modification as discussed above, and the modification occurs in place in the entry database 710 without removing the matching multi-entry. For example, if template (i) in the multi-template 1612 is non-null, the entry (i) in the multi-entry being altered is modified according to template (i) in multi-template 1612 using single entry modification (step 1708 of Fig. 17, where template (i) in multi-template 1612 is used as template 2). The values of the non-null fields of the template (i) are used to replace the values in the corresponding fields of the existing entry being modified.

15 For every null template in the given multi-template 1612, the corresponding entry in the matching multi-entry is deleted (*i.e.*, for every template (i) in the multi-template 1612 that is null, entry (i) in the multi-entry to be modified is deleted) (step 1910). If there are more multi-entries in the database that match the multi-template 1610, these steps are repeated to modify the remaining matching multi-entries in the entry database 710 (step 1912). When the search engine 708 completes the modifications, it returns a number representing the number of multi-entries that were modified (step 1914).

20 An exception case arises, however, when dealing with the multi-entry modify/delete operation. In multi-entry/multi-template matching, because an entry in a multi-entry may be used to match more than one template in a multi-template, a multi-entry may match a multi-template and have less entries than the multi-template has templates. If such a multi-entry matches the multi-template 1610 and is to be modified by the modify/delete operation consistent with the present invention, more than one template in the multi-template 1612 will be used to modify or delete a single entry in the multi-entry to be modified. In this case, if any such template in the multi-template 1612 (*i.e.*, templates corresponding to templates of the multi-template 1610 that match the entry of the multi-entry to be modified) indicates a delete operation

25

30

(i.e., is null), then the entry in the multi-entry to be modified is deleted. Otherwise, if all such templates in the multi-template 1612 indicate modify operations, then all of the templates are used for modification, but the decision on the order in which to apply the modifications will be implementation specific and does not have any significant bearing on the modify/delete operation as a whole.

(4) Modification Notification

Regarding in-place modifications, in another implementation consistent with the present invention, a user can register interest in entries or multi-entries in the database and receive notification back to the user or an event catcher as before. In this implementation, when a user registers interest in entries or multi-entries in a database by supplying a template or multi-template, the search engine 708 does not monitor incoming entries or multi-entries as previously done. When implementing notification of in-place modifications, it monitors entries or multi-entries that have been changed or modified in the database. The monitoring of modified entries or multi-entries in the database as opposed to incoming entries to the database is the main difference between the previous notification operation and the present notification operation relating to in-place modification. However, many of the basic aspects of the notification remain the same.

Generally, a template 1614 or multi-template 1616 is used to register interest in changing entries and multi-entries. Upon the changing of a multi-entry or entry during an in-place modification, the search engine 708 checks to determine if the changed entry or multi-entry now matches the registered template or multi-template, or if the changed multi-entry or entry transitioned from matching the template or multi-template to not matching the template or multi-template. Hence, the notification operation searches for transitions in matching status, either from matching to not matching or not matching to matching.

Turning next to Fig. 20, a flowchart is shown which illustrates one implementation of a methodology implemented by search engine 708 to process a notify request received from a user to monitor changes in modified entries or multi-entries. In this case, the notify request may be issued by the user to register interest in specified entries or multi-entries to determine whether they change matching status in relation to the supplied template 1614 or multi-template 1616. When an entry or multi-entry is modified, its matching status can transition by changing from

matching the template 1614 or multi-template 1616 to not matching it, or by changing from not matching to matching. When an entry or multi-entry changes its matching status with regard to a template 1614 or multi-template 1616 in a notification operation request, the database manager 704 notifies either the user or an event catcher (see Fig. 2), as desired. When a user invokes a notify request, a lease time is specified which indicates how long the database manager 704 will register the notification. After the specified time has elapsed, the notify registration is canceled.

Thus, as depicted in Fig. 20, when database manager 704 receives a new notify request (step 2002), entries or multi-entries are monitored when they are modified until the registration time expires (steps 2004 and 2006). If an entry or multi-entry is modified in the entry database 710, the search engine 708 determines whether the entry or multi-entry changed its matching status with regard to the template 1614 or multi-template 1616 received in the notification request (step 2008 and 2010). The search engine 708 determines whether the modified entry or multi-entry went from matching the template or multi-template in the request to not matching it or vice versa. If a transition of the matching status occurs, the user or the event catcher is notified (step 2012). Additional entries or multi-entries that have been modified in place continue to be monitored as long as the registration time has not expired (step 2004).

Type-Safe Attribute Matching

Another implementation consistent with the present invention involves attributes that represent various values of database entries represented as objects. In this implementation, by associating attributes to the fields of entries and using entries, multi-entries and their operations, type-safe attribute matching can be achieved. Type-safe attribute matching involves the searching of different types of attributes in a database and allows attributes to be strongly typed. It further allows the enforcing of relationships between attributes and allows efficient referencing and manipulation of attributes by type.

Using type-safe attribute matching, attributes and names are strictly associated with each other (*i.e.*, a pages per minute attribute is associated only with a printer) as well as attributes and their values (*i.e.*, pages per minute is associated only with an integer). Type-safe attribute matching also allows enforcement of relationships between pairs of attributes and their values. It also enables efficient searching for objects by type while allowing objects of many different types to be stored in a single database. Additionally, it permits modifications that affect only a

specified type of object in the database, although the database may store many different types of objects.

Consistent with the present invention, type-safe attribute matching is achieved by associating attributes with the fields of an entry. The value stored in the field of the entry is the value of the associated attribute. Since a set of attributes, an entry in this case, can be an instance of a Java™ class, the class provides strong typing of both the set and the individual attributes. For instance, the following example illustrates an entry representing a typical printer with fields representing the attributes of the printer (*i.e.* pages per minute, resolution, etc.), as well as entries representing more generic information such as names, product data, and physical location:

```
public class Printer implement Entry {
    Integer ppm;      //    pages per minute
    Integer dpi;      //    resolution in dots per inch
    Boolean duplex;   //    supports two-sided printing
    Boolean color;    //    color or black-only
}

public class Name implements Entry {
    String name;      //    the user-friendly name of the service
    String description; //    free-form description of the service
    String language;   //    language (e.g., English, French) used in
                        //    the
                        //    above
}

public class Type implements Entry {
    String type;      //    the general type of service
    String vendor;    //    vendor of product that implements the
service
    String model;     //    model number/name of product
    String version;   //    version number of product
}

public class Location implements Entry {
    Integer floor;    //    what floor the service is on
    String building;  //    what building it's in
    String room;      //    what room it's in
}
```

As shown, there is strict typing of fields as well as the entries, and these fields and their types will always be associated with these entries and their types. For example, "ppm" or "pages per minute" is of type "Integer".

Further description of an item may be achieved using a multi-entry. Using a multi-entry, a typical printer may also be described as a set of sets of attributes. When describing a typical printer in this manner, additional aspects may be included such as a location, name or machine type, each aspect represented by another entry in the multi-entry. The set of entries in the example above form an exemplary multi-entry representing the attributes of a typical printer.

The entry and multi-entry operations may be used to search and manipulate these type-safe sets of attributes. These operations include take, write, notify, read, entry modify, multi-entry add, multi-entry modify/delete, and modification notification. It should be noted that other known database operations may be performed in the database systems consistent with the present invention.

For exemplary purposes, assume the existence of two different printers described by the following multi-entries stored in a database. Each printer may have certain properties described by the attributes associated with the fields of the entries in the multi-entry. In this hypothetical example, the first printer has an English name "quick" and a German name "schnell," both with associated descriptions. As shown in the example, it is located on floor 1 of building ABC in room 1202 and also has associated printing characteristics (e.g. ppm=10, dpi =600, etc.) Such a printer may be represented by the following multi-entry:

Multi-Entry 1 =

{Name	[name	=	"quick",
	description	=	"yadal",
	language	=	"English"],
Name	[name	=	"schnell",
	description	=	"yahdal",
	language	=	"German"],
Location	[floor	=	1,
	building	=	"ABC",
	room	=	"1202"],
Printer	[ppm	=	10,
	dpi	=	600,

30

```

duplex      = true,
color       = false}}

```

5 A second printer, having an English name "slow" and located on floor 2 of building XYZ in room 101, also has associated printing characteristics (e.g. ppm=3, dpi=300, etc.). The second printer may be represented by a multi-entry as follows:

```

10 Multi-Entry 2 = {Name      {name      = "slow",
                             description = "yada2",
                             language   = "English"},
                             Location    {floor   = 2,
                             building    = "XYZ",
                             room        = "101"},
15                             Printer    {ppm     = 3,
                             dpi         = 300,
                             duplex      = false,
                             color      = true}}

```

20 Now, suppose a search is required that returns all printers with certain attributes represented by the fields of entries. For example, if a user sought a printer represented in the database that was located in building "XYZ" and was a color printer, the multi-template in the search might be represented as follows:

```

25 Multi-Template 1 = {Location {floor   = null,
                             building  = "XYZ",
                             room      = null},
                             Printer   {ppm     = null,
                             dpi        = null,
30                             duplex    = null,
                             color     = true}}

```

35 A user might call a read operation supplying Multi-Template 1 as the parameter. If applied to the database containing Multi-Entries 1 and 2, the Multi-Entry 2 representing the second printer from the previous example would be returned because the Multi-Template 1, shown here, matches Multi-Entry 2. The search engine 708 searches the database, bypassing items of


```
color = null))
```

The search engine 708 locates all multi-entries in the database matching Multi-template 2 and modifies them according to the second parameter in the operation, as discussed in the multi-entry modify/delete operation method. Multi-Entry 1 from the previous example matches the Multi-Template 2 used in the operation parameter, and therefore, Multi-Entry 1 is changed according to the second parameter of the operation. This operation deletes the location attributes, replaces the dpi attribute with a value of 100 and yields a modified Multi-entry 1:

Multi-Entry 1 =

```
{Name      {name      = "quick",
            description = "yadal",
            language    = "English"},
  Name      {name      = "schnell",
            description = "yahdal",
            language    = "German"},
  Printer   {ppm       = 10,
            dpi        = 100,
            duplex      = true,
            color       = false}}
```

Hence the search engine sought all items in the database possessing attributes of a requested type and modified them in a type-safe manner, thus implementing type-safe attribute matching. It should be noted that the examples shown are for purposes of illustration only.

Numerous variations and modifications will become apparent to those skilled in the art once the above disclosure is fully appreciated. It is intended that the following claims be interpreted to embrace all such variations and modifications.

WHAT IS CLAIMED IS:

1. A method for searching items in a database, the database storing a plurality of items, each item comprising a plurality of entries, the method comprising:

receiving a request including a multi-template comprised of a plurality of templates;

5 comparing the multi-template to an item in the database to determine whether the item matches the multi-template; and

returning the item based upon a result of the comparison.

2. The method of claim 1, wherein the comparing step comprises the step of:

comparing each template in the multi-template with the entries in the item to determine if each template has a matching entry, each entry being allowed to match more than one template.

3. The method of claim 1, further including the step of:

repeating the comparing step until a match is found or all items in the database are compared.

4. The method of claim 3, wherein the returning step includes the step of:

returning a null value if no matches are found.

5. The method of claim 1, wherein the receiving step includes the step of:

receiving a request for more than one item.

6. The method of claim 5, wherein the returning step includes the step of:

returning more than one item.

7. A method for notification of the arrival of a specified item in a database comprising the steps of:

receiving a request to monitor the database for the arrival of a specified item, the request including a multi-template;

determining, upon the arrival of an item, whether the item matches the multi-template in the request; and

5 returning notification based on the determination.

8. The method of claim 7, wherein the receiving step includes the step of:
receiving a request including a lease time indicating the length of time the database is to monitor incoming items.

9. A method for searching items in a database, the database storing a plurality of items, each item comprising a plurality of entries, the method comprising:

providing a request including a multi-template comprised of a plurality of templates; and

5 receiving an item matching the multi-template based on a comparison of the multi-template to entries in the database.

10. The method of claim 9, wherein the request includes a take or read operation.

11. The method of claim 9, wherein the receiving step includes the step of:
receiving a null value if the database contains no items matching the multi-template.

12. A method for searching items in a database, the database storing a plurality of items, each item comprising a plurality of entries, the method comprising:

providing a request including a take or read operation and a multi-template comprised of a plurality of templates;

5 selecting an item in the database to compare to the multi-template;

comparing the multi-template to one of the items in the database by comparing each template in the multi-template with the entries in the one of the items to determine if each template has a matching entry, each entry being allowed to match more than one template;

returning the one of the items or a null value based on a result of the comparison.

13. The method of claim 12, further comprising:

repeating the selecting and comparing steps until a match is found or all items in the database have been compared.

14. A data processing system for searching for items in a database, the system comprising:

a CPU;

a memory coupled to the CPU, the memory storing a database comprising a plurality of items, each item comprising a plurality of entries;

wherein the memory stores a search program which is executable by the CPU to receive a request including a multi-template comprised of a plurality of templates, compare the multi-template to an item in the database to determine whether the item matches the multi-template, and

return the item based upon a result of the comparison.

15. A data processing apparatus having a database storing a plurality of items comprising a plurality of entries, the apparatus comprising:

a receiving component configured to receive a request including a multi-template comprised of a plurality of templates;

5 a comparing component configured to compare the multi-template to an item in the database to determine whether the item matches the multi-template; and

a returning component configured to return the item based upon a result of the comparison.

16. The apparatus of claim 15, wherein the comparing component comprises:

a comparing component configured to compare each template in the multi-template with the entries in the item to determine if each template has a matching entry, each entry being allowed to match more than one template.

17. The apparatus of claim 15, further including:

a repeating component configured to repeat the comparing step until a match is found or all items in the database are compared.

18. A data processing apparatus having a database storing a plurality of items comprising a plurality of entries, the apparatus comprising:

a providing component configured to provide a request including a multi-template comprised of a plurality of templates; and

5 a receiving component configured to receive an item matching the multi-template based on a comparison of the multi-template to entries in the database.

19. A computer program product comprising a computer readable medium having computer readable code embodied therein for processing data in a database by:

receiving a request including a multi-template comprised of a plurality of templates;

5 comparing the multi-template to an item in the database to determine whether the item matches the multi-template; and

returning the item based upon a result of the comparison.

20. The product of claim 19, wherein the comparing further comprises:

comparing each template in the multi-template with the entries in the item to determine if each template has a matching entry, each entry being allowed to match more than one template.

21. The product of claim 19, further comprising:

repeating the comparing step until a match is found or all items in the database are compared.

22. A computer program product comprising a computer readable medium having computer readable code embodied therein for processing data in a database by:

providing a request including a multi-template comprised of a plurality of templates; and

5 receiving an item matching the multi-template based on a comparison of the multi-template to entries in the database.

23. A computer program product comprising a computer readable medium having computer readable code embodied therein for processing data in a database, the computer readable medium comprising:

means for receiving a request including a multi-template comprised of a plurality of templates;

means for comparing the multi-template to an item in the database to determine whether the item matches the multi-template; and

means for returning the item based upon a result of the comparison.

1/18


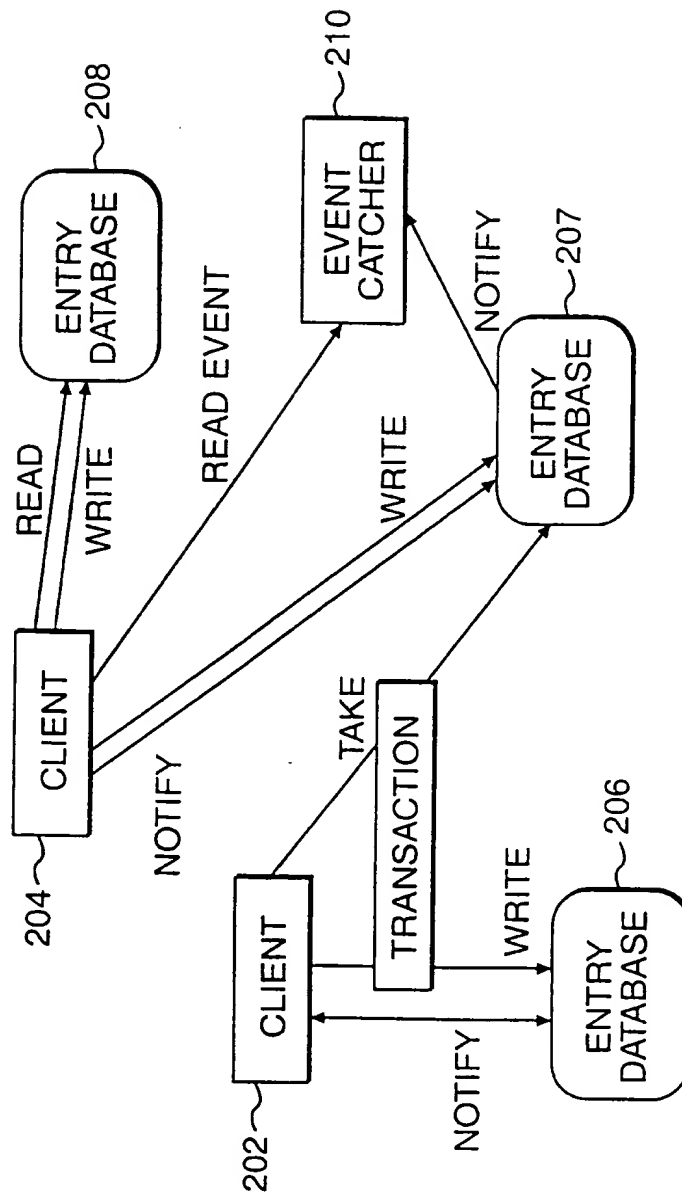


FIG. 1

2/18

**FIG. 2**

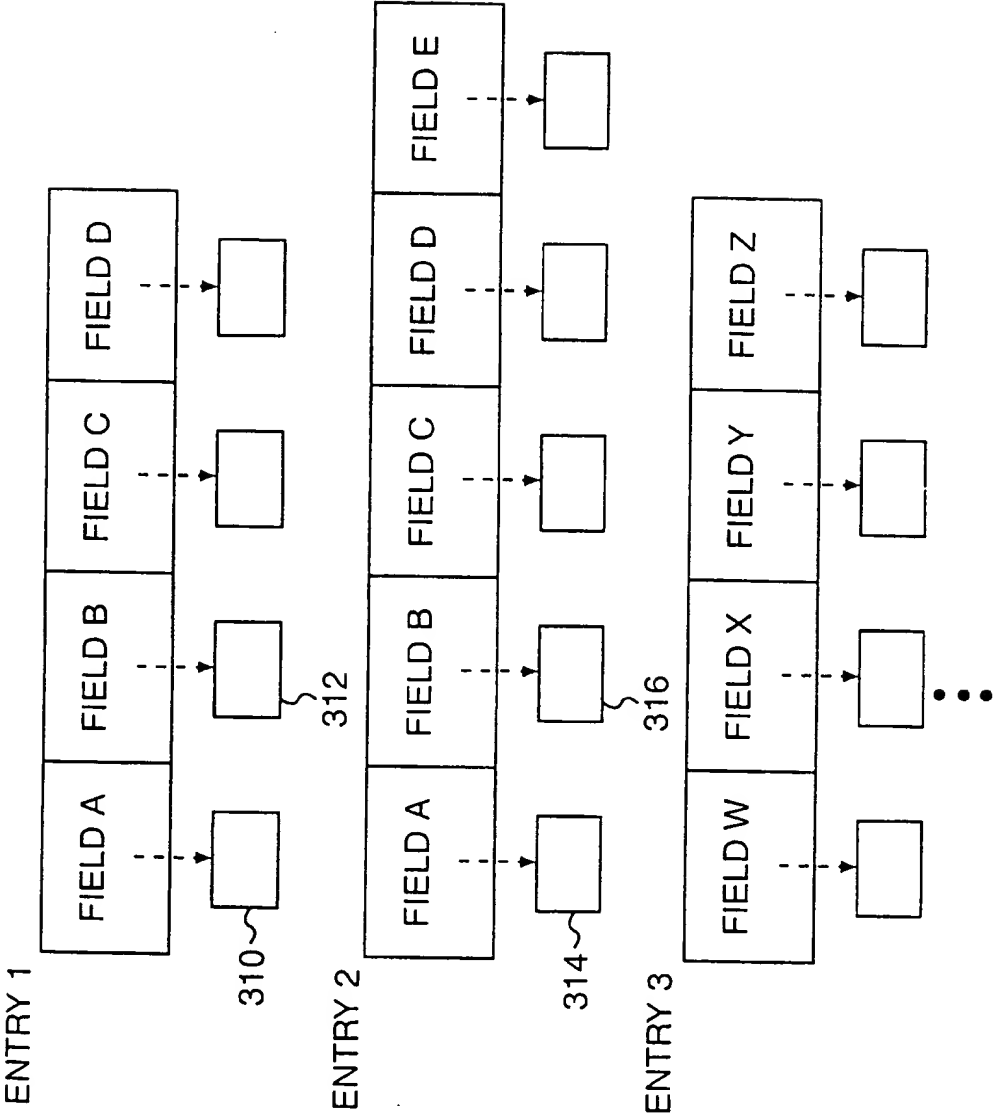


FIG. 3


4/18

WRITE: [ENTRY]

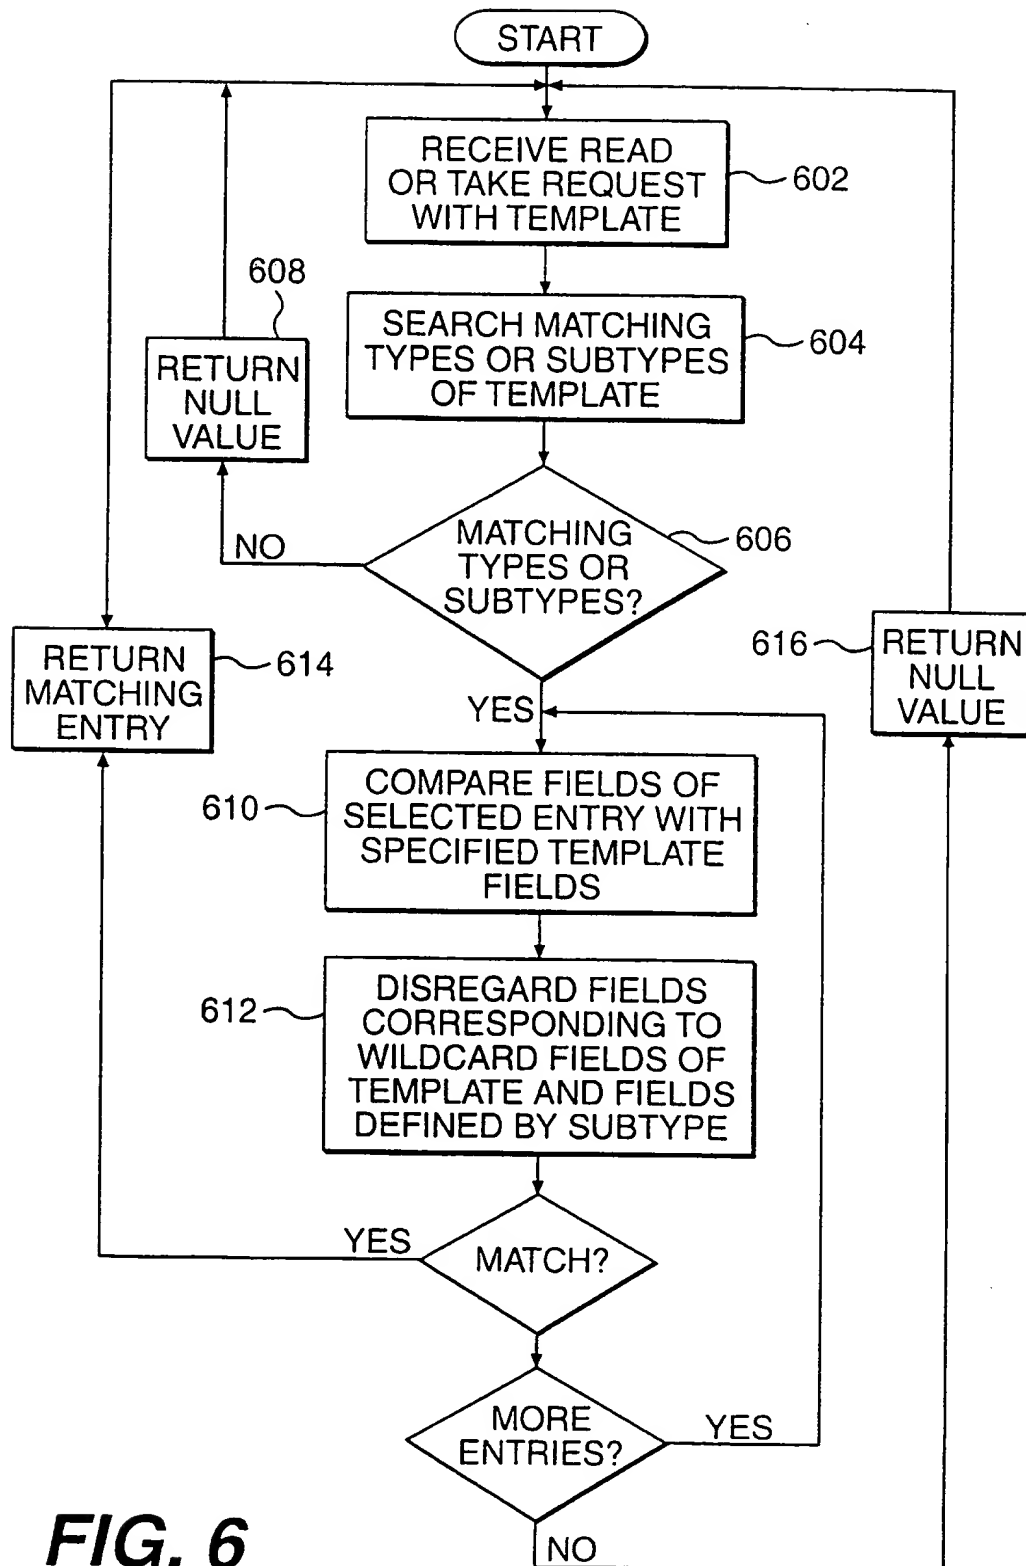
READ: [TEMPLATE]

TAKE: [TEMPLATE]

NOTIFY: [TEMPLATE]

FIG. 4**FIG. 5**

5/18

**FIG. 6**

6/18


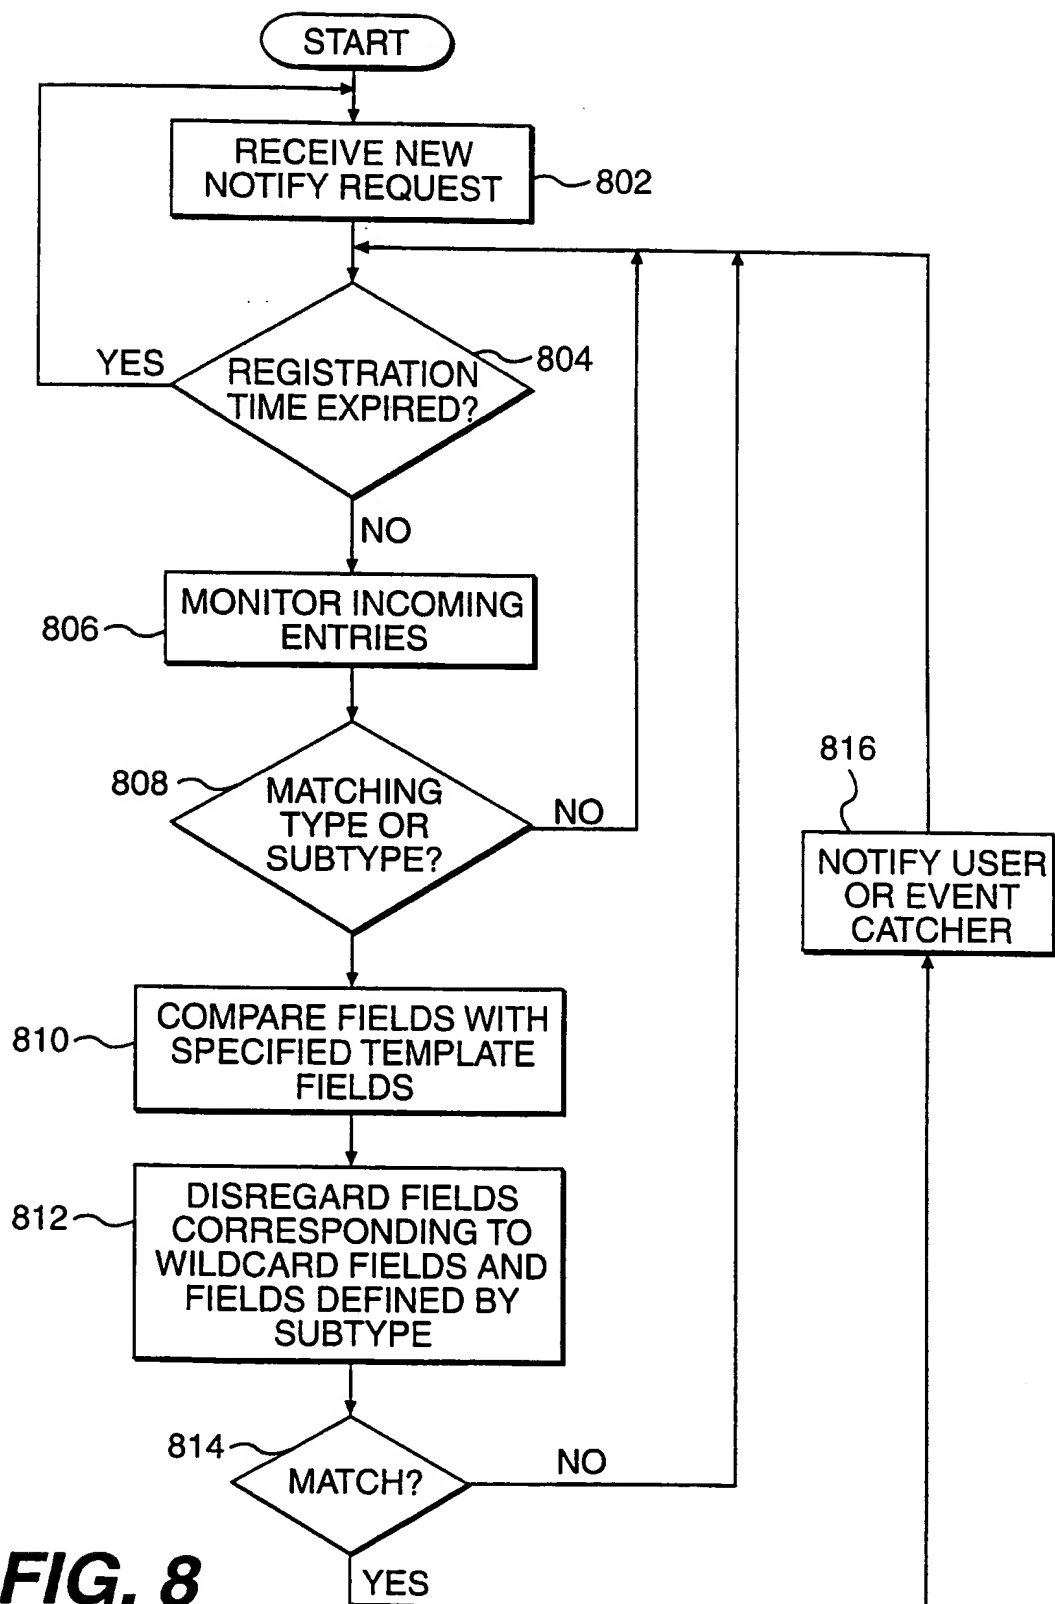
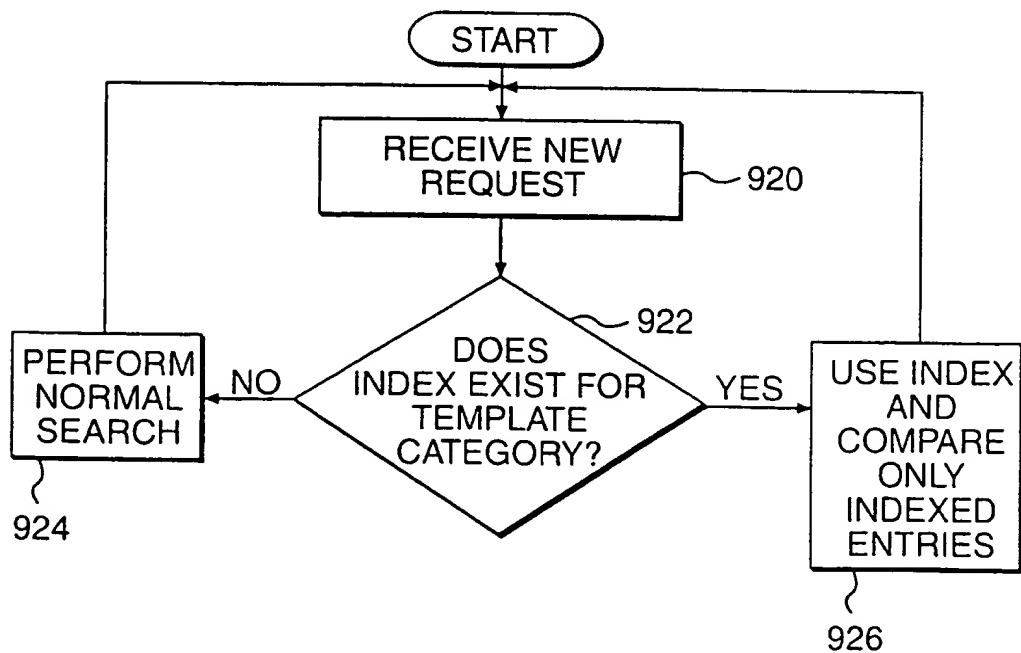



FIG. 7

7/18

**FIG. 8**

8/18



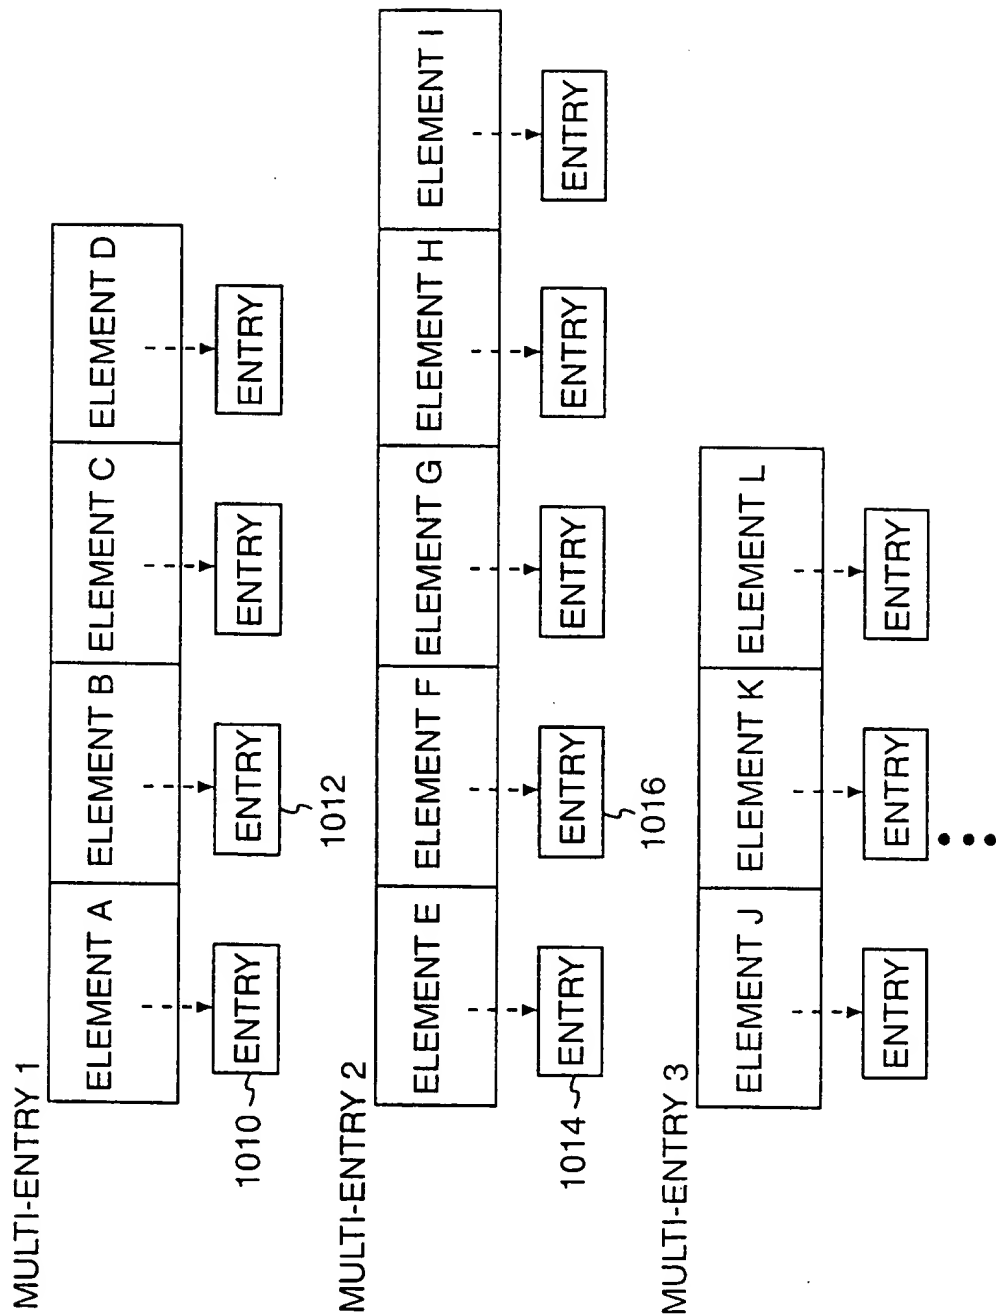


FIG. 10

10/18

WRITE: [MULTI-ENTRY]

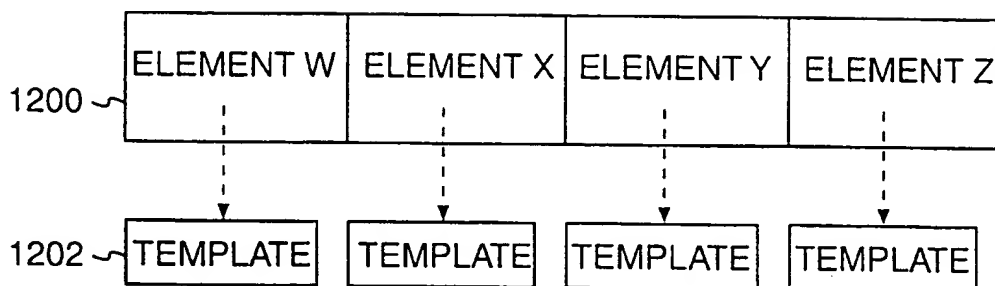
READ: [MULTI-TEMPLATE]

TAKE: [MULTI-TEMPLATE]

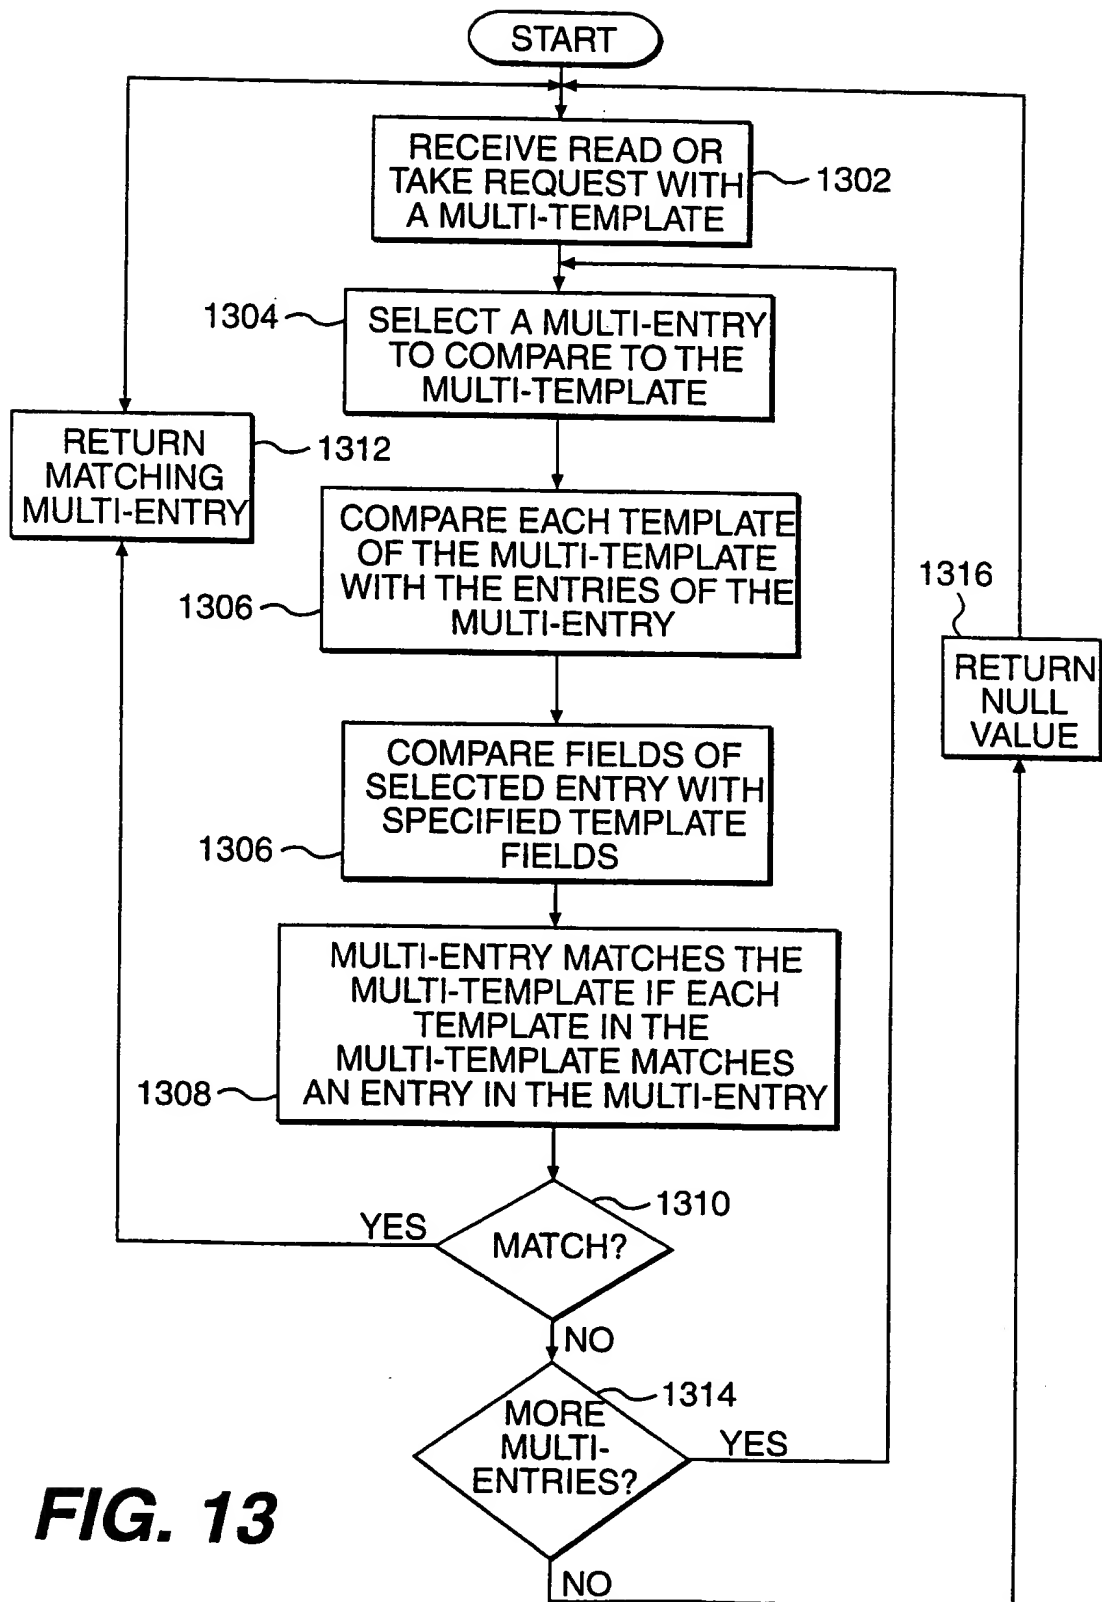
NOTIFY: [MULTI-TEMPLATE]

FIG. 11

MULTI-TEMPLATE

**FIG. 12**

11/18

**FIG. 13**


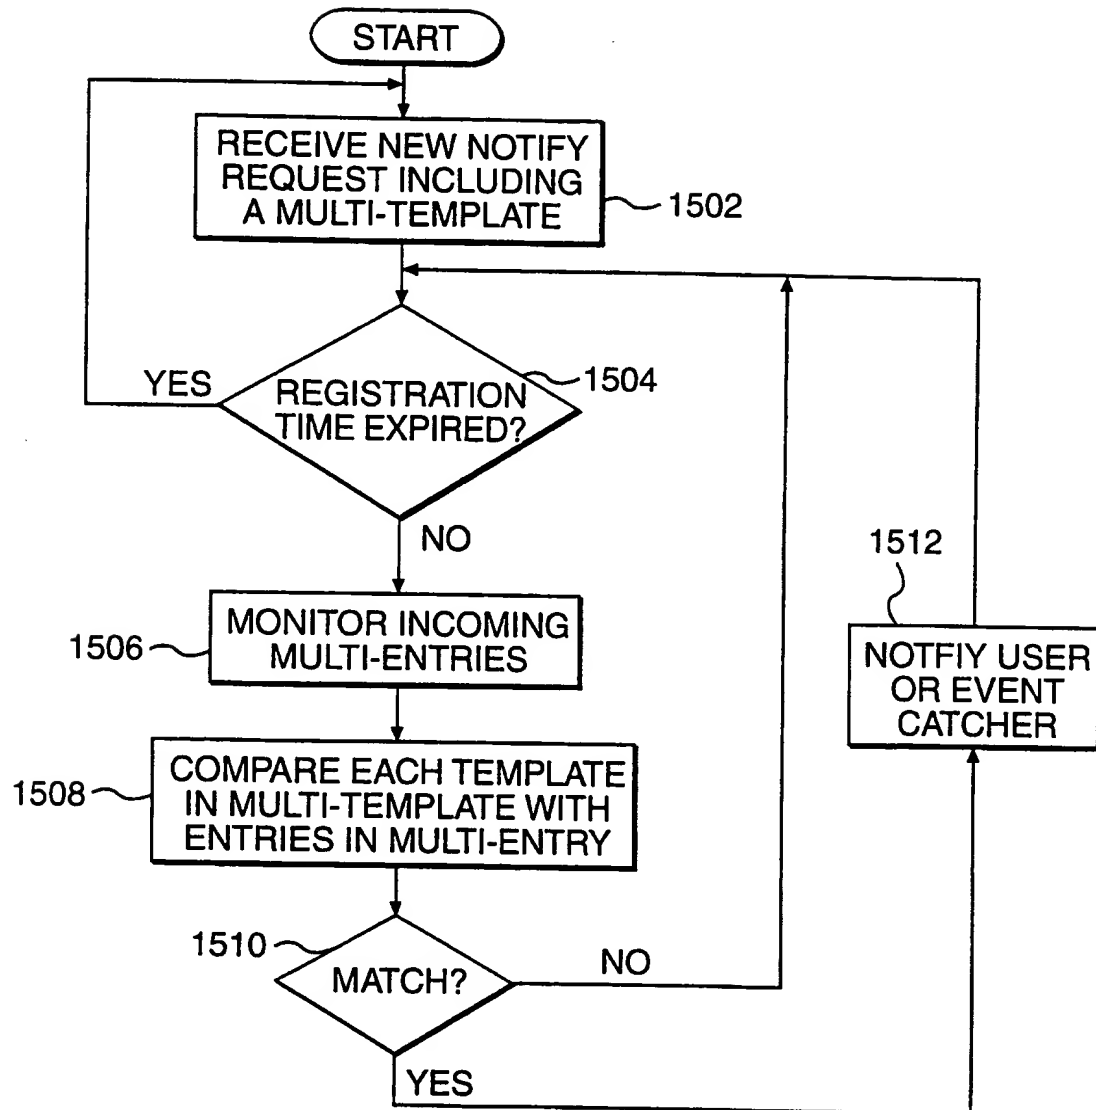


FIG. 14

13/18

**FIG. 15**

14/18


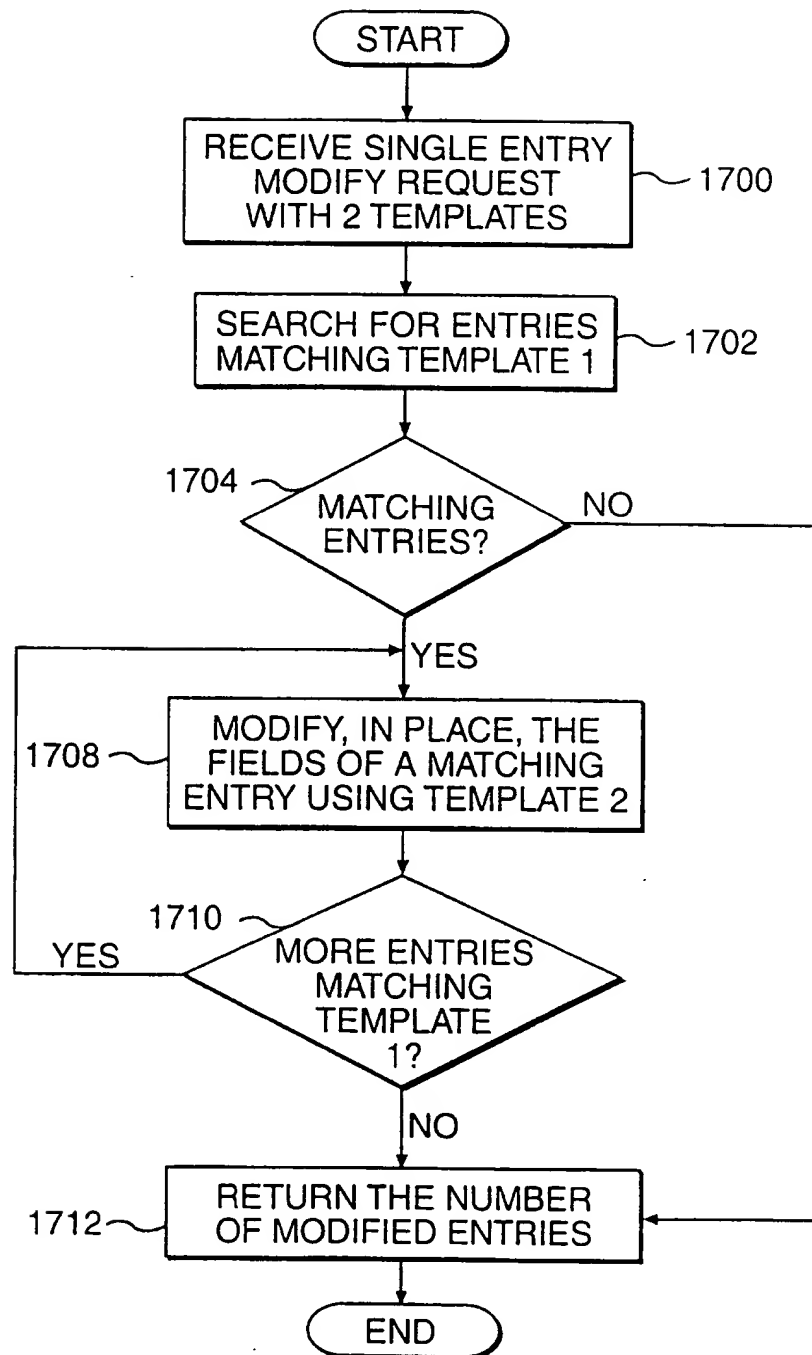
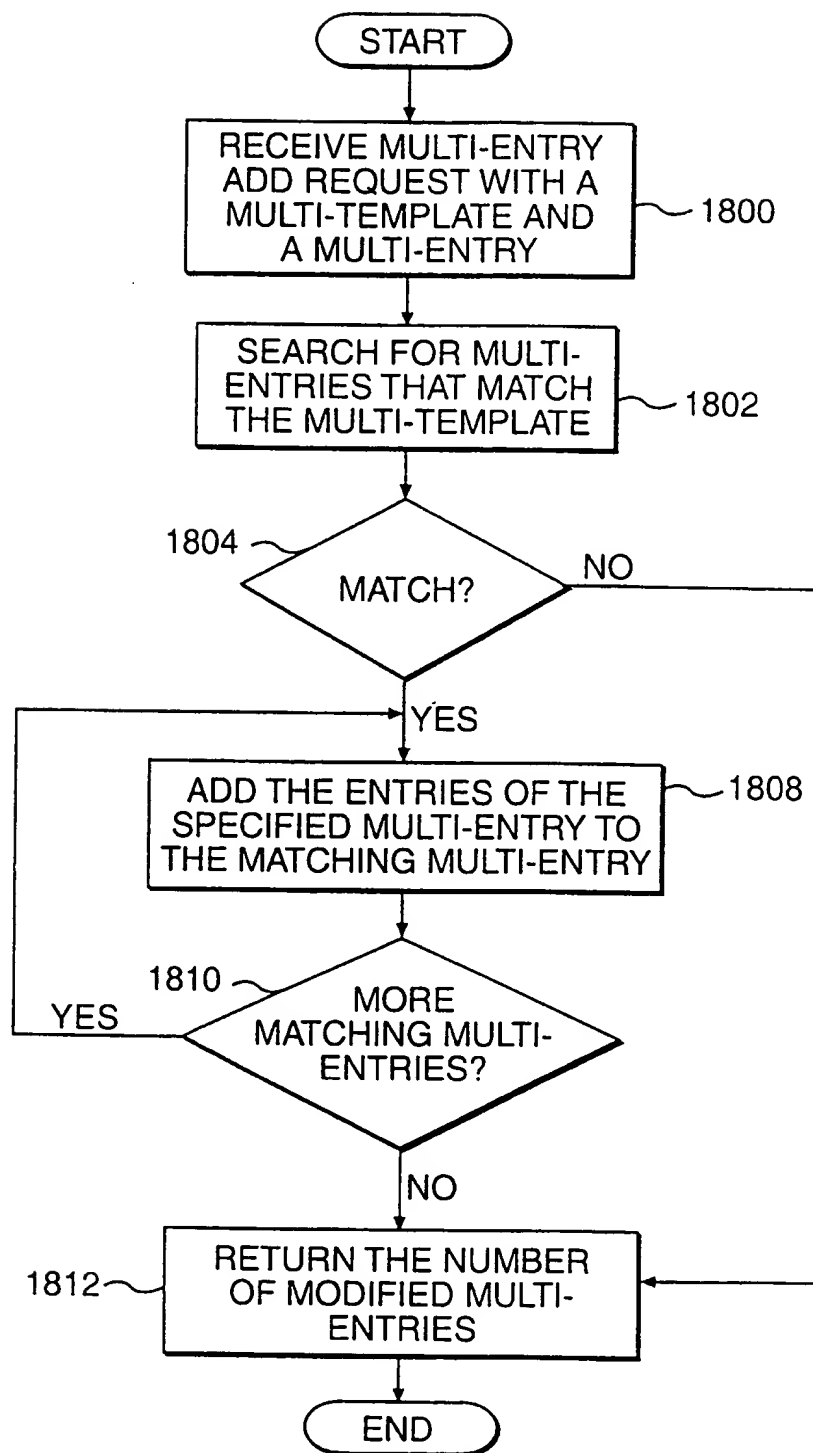


FIG. 16

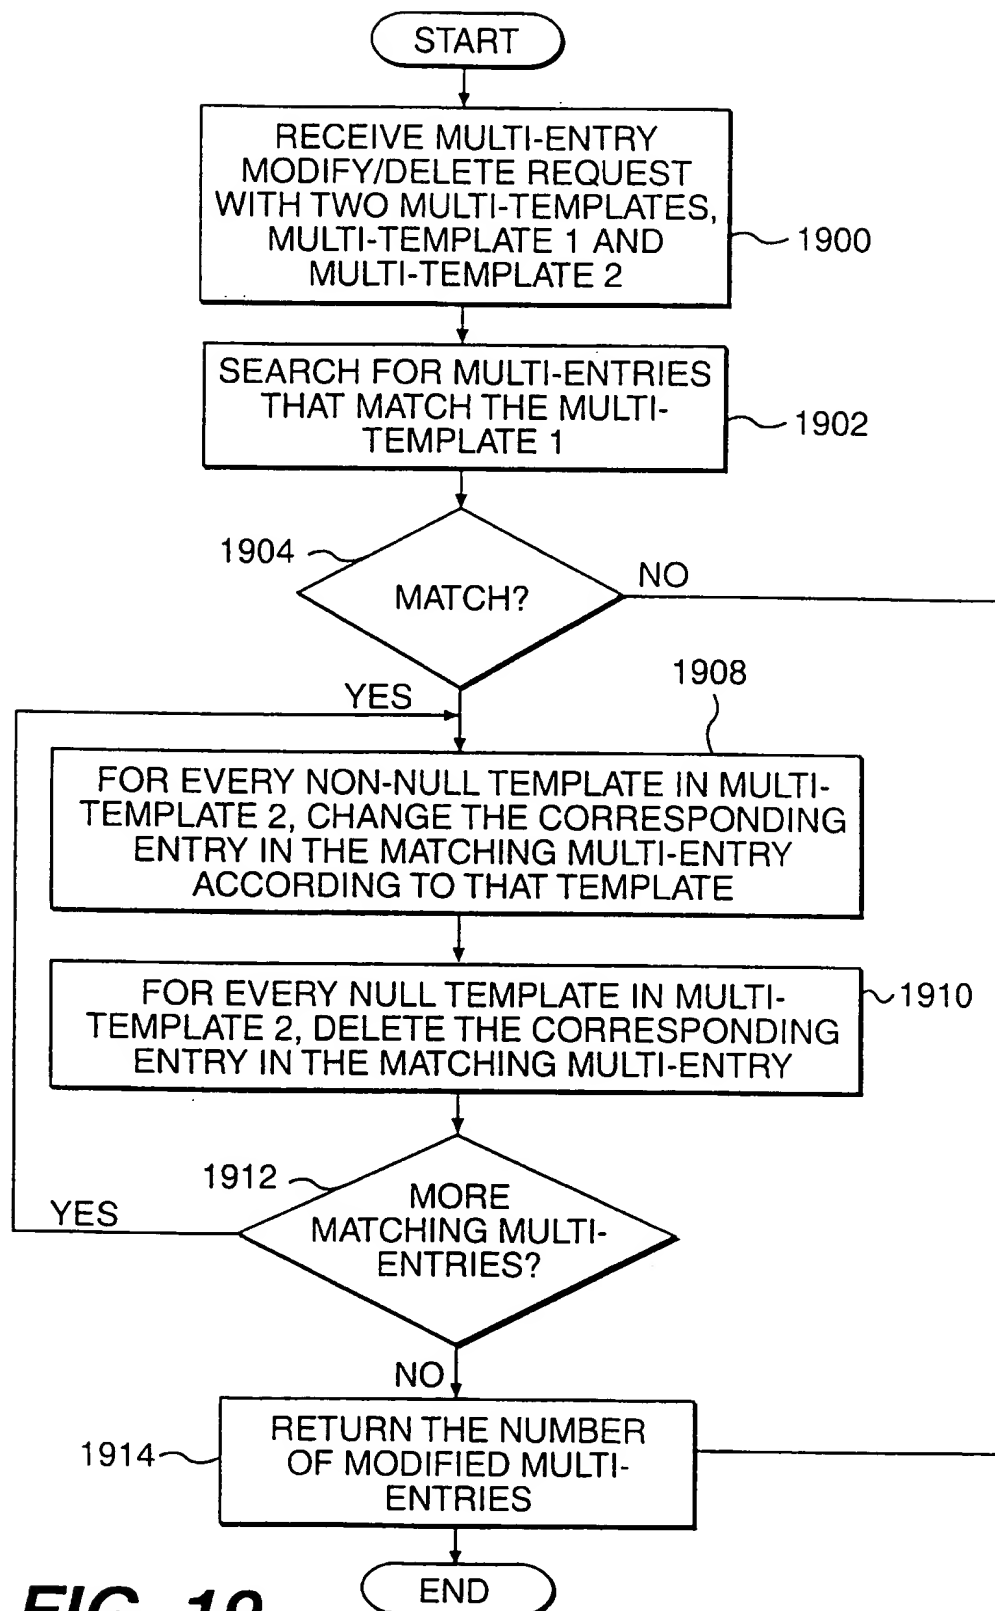
15/18

**FIG. 17**

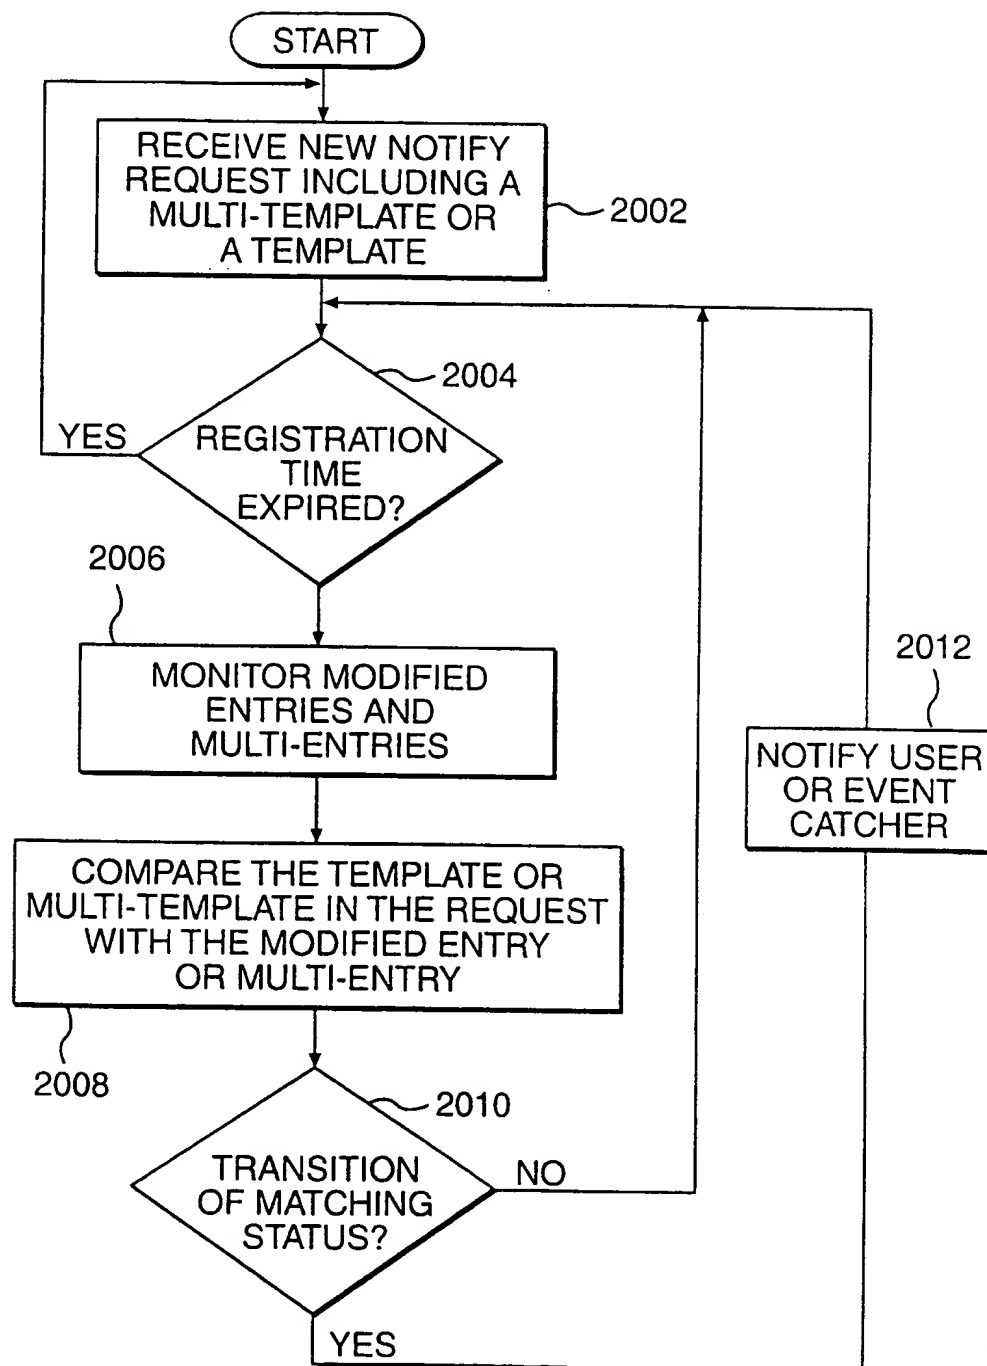
16/18

**FIG. 18**

17/18

**FIG. 19**

18/18

**FIG. 20**

INTERNATIONAL SEARCH REPORT

Intern. al Application No

PCT/US 99/04146

A. CLASSIFICATION OF SUBJECT MATTER
IPC 6 G06F17/30

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 6 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	GOTTLOB G ET AL: "EXTENDING OBJECT-ORIENTED SYSTEMS WITH ROLES" ACM TRANSACTIONS ON INFORMATION SYSTEMS, vol. 14, no. 3, 1 July 1996 (1996-07-01), pages 268-296, XP000595964 ISSN: 1046-8188 page 270, line 31 - page 271, line 8	1,7,9, 12,14, 15,18, 19,22,23
A	BERTINO E ET AL: "OBJECT-ORIENTED DATABASE MANAGEMENT SYSTEMS: CONCEPTS AND ISSUES" COMPUTER, vol. 24, no. 4, 1 April 1991 (1991-04-01), pages 33-47, XP000227753 ISSN: 0018-9162 page 34, column 2, line 18 - page 34, column 2, line 22 ----- -/-	1,7,9, 12,14, 15,18, 19,22,23

☒ Further documents are listed in the continuation of box C.

☐ Patent family members are listed in annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search

30 July 1999

Date of mailing of the international search report

05/08/1999

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3018

Authorized officer

Katerbau, R

INTERNATIONAL SEARCH REPORT

Intern 1al Application No

PCT/US 99/04146

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	<p>BEECH D ET AL: "OBJECT DATABASES AS GENERALIZATIONS OF RELATIONAL DATABASES" COMPUTER STANDARDS AND INTERFACES, vol. 13, no. 1 / 03, 1 January 1991 (1991-01-01), pages 221-230, XP000288495 ISSN: 0920-5489 page 222, column 1, line 14 - page 226, column 1, line 11</p> <p>-----</p>	<p>1,7,9, 12,14, 15,18, 19,22,23</p>